JPL D-18441

Report on the Loss of the Mars Climate Orbiter Mission

JPL Special Review Board

November 11, 1999

Jet Propulsion Laboratory
California Institute of Technology

54 pp

D 18441

Table of Contents

D 18441

Table of Figures

Signature Page

Richard Brace (JPL, Office 514)

John Casani, Chair (JPL, retired)

Robert Farquhar (APL)

Norm Haynes (JPL, Section 400)

Frank Jordan (JPL, Section 704)

Charles Kohlhase (JPL, retired)

Bob Mitchell (JPL, Section 990)

Robert J. Polutchko (LMC, retired)

Al Schallenmuller (LMA, retired)

John Slonski (JPL, Section 313)

Robert H. Tolson (GWU)

Acronym List

0 18441

ACG	Attitude Control Group
AMD	Angular Momentum Desaturation
APL	Applied Physics Laboratory
DSN	Deep Space Network
DTR	Detailed Technical Review
EME	Earth Mean Equator (and equinox)
FR	Failure Report
GWU	George Washington University
GDS	Ground Data System
I/O	Input/Output
IBIT	Impulse Bit
ISA	Incident/Surprise/Anomaly
JPL	Jet Propulsion Laboratory
Km	Kilometer
LMA	Lockheed Martin Astronautics
LMC	Lockheed Martin Corporation
LOS	Line Of Sight
MCO	Mars Climate Orbiter
MGS	Mars Global Surveyor
MME	Mars Mean Equator (and equinox)
MO	Mars Observer
MOI	Mars Orbit Insertion (MOI)
MPL	Mars Polar Lander
MSOP	Mars Surveyor Operations Project
OD	Orbit Determination
ODP	Orbit Determination Program
ORR	Operational Readiness Reviews
ORT	Operational Readiness Test
P/FR	Problem/Failure Report
P1	Periapsis Altitude on Orbit One
P2	Periapsis Altitude on Orbit Two
PEM	Project Element Manager
PM	Program Manager
PRM	Periapsis Raise Maneuver
SIS	Software Interface Specification
SMDP	Software Management and Development Plan
SPAS	Spacecraft Performance and Analysis Software
SRDS	Software Requirements and Design Specification
STL	Spacecraft Test Lab
TCM	Trajectory Correction Maneuvers

1 Executive Summary

18441

The Mars Climate Orbiter (MCO) was launched on December 11, 1998. The MCO was to arrive at Mars and begin orbit insertion on September 23, 1999. The Mars Orbit Insertion (MOI) burn, a 16-minute maneuver to slow the spacecraft and enable capture into an orbit around Mars, began on schedule. Five minutes into the maneuver, and approximately 49 seconds before the anticipated time for loss of communication, the MCO was occulted by Mars. Thereafter, no contact with the spacecraft could be established.

On September 24, 1999, an internal JPL team (the MCO Peer Review Team) was appointed to help investigate the reason for the loss of spacecraft signal. The Peer Review Team's findings are summarized in section 2, Peer Review Team Investigation, and detailed in Appendix 1.

Also on September 24, 1999, JPL appointed a Special Review Board (the Board), which included members from JPL, industry and academia, as follows:

- Richard Brace (JPL, Office 514)
- Robert Farquhar (APL)
- Frank Jordan (JPL, Section 704)
- Bob Mitchell (JPL, Section 990)
- Al Schallenmuller (LMA, retired)
- Robert H. Tolson (GWU)
- John Casani, Chair (JPL, retired)
- Norm Haynes (JPL, Section 400)
- Charles Kohlhase (JPL, retired)
- Robert J. Polutchko (LMC, retired)
- John Slonski (JPL, Section 313)

The Board was tasked to:

1) Determine how the mission loss occurred.

2) Provide an analysis of why the established processes failed.

3) Recommend ways to strengthen or augment established processes to help prevent future mishaps.

With regard to task (1) above, the Peer Review Team and the Board determined that the mission loss was precipitated by an error in the software program that generated the Angular Momentum Desaturation (AMD) files. This error was revealed when, on September 29, 1999, Lockheed-Martin Astronautics (LMA) reported that the files containing the magnitudes of the small forces impulses applied to the spacecraft had been delivered in English units (pounds-force seconds) instead of metric units (Newton-seconds). The interface agreements specify metric units. (Appendix 2 is a simplified technical discussion of the orbit-determination process and its dependence on accurate knowledge of the forces acting on a spacecraft.)

The discrepancy in the AMD files led to an underestimate by a factor of 4.5 of the influence of small forces on the spacecraft trajectory. Over the 9-month mission, the cumulative effect of these twice-a-day impulses would have produced an error at Mars of more than 10,000 km. Most of this error was removed by the four trajectory correction maneuvers (TCMs) performed during cruise to Mars. After the fourth and final TCM, performed on September 15, 1999, the remaining effect of the small forces error was 169 km. Post-failure analysis has shown that the spacecraft's minimum altitude would have been 57 km. The 169-km difference between this prediction and the actual first periapsis altitude represents a major discrepancy in the mission's navigation.

Although the software error and the resulting navigation inaccuracies were the initial factors, the following other factors contributed to the mission loss:

- The failure to detect the units error, either during development or in-flight operations.
- The failure to properly compensate for observed in-flight discrepancies.
- The failure to follow the baseline mission-risk strategy.

With regard to task (2) above, the Board finds that there were numerous variances from established processes. Conscientious application of those processes would have:

(a) Revealed the error prior to launch, via software reviews.

(b) Revealed the error during cruise, through the use of navigation operational guidelines.

(c) Mitigated the effects of the error, through project risk-management processes.

Relative to task (3) above, the Board has findings and recommendations in 13 specific areas. Section 4 contains these findings and recommendations. Major areas include the Software Development Process and Operational Readiness Testing, Project Management and Decision Making, Navigation Process and Relationship to the Project, Effectiveness of Reviews, Operational Anomaly Reporting and Tracking, Mission Assurance, and Flight Operations Training. Recommendations relevant to the Mars Polar Lander in the short term are indicated by "(MPL)" following the recommendation.

2 Peer Review Team Investigation

As is noted above, contact with the MCO was lost on September 23, 1999. On September 24, 1999, JPL appointed the MCO Peer Review Team, chaired by Dr. Frank Jordan, to help the MCO team investigate the loss of the MCO signal.

The Peer Review Team performed a fault-tree analysis to isolate the possible cause for this signal loss. The analysis began with the investigation of the following possible causes:

- Failure of the spacecraft while it was occulted from view.
- Failure of the Deep Space Network (DSN).
- The possibility that the spacecraft entered the Martian atmosphere and was destroyed.

Through investigation and elimination, the Peer Review Team reached the following conclusions:

- The loss of mission occurred when the MCO entered the Martian atmosphere (see the Note below).
- The MCO entered the Martian atmosphere because the spacecraft approached at an incorrect trajectory.
- The trajectory error was caused by errors in small forces modeling.
- The modeling error occurred because small forces information was delivered in pounds-force seconds, instead of the required Newton-seconds.

Figure 1 is a truncated fault tree of the possible causes of the loss of signal. The fault tree is truncated to exclude root error sources for those areas that the team had ruled out as possible contributors to the mishap.

Appendix 1 is a detailed discussion of the Peer Review Team's investigation and conclusions. This appendix begins with a brief description of the Mars Orbit Insertion (MOI) maneuver, which preceded the loss of the MCO signal. Following that description, the appendix traces possible reasons for the loss of signal, providing supporting arguments for the Team's conclusions and the bases for ruling out other possible explanations.

Note:

Approximately half way through the MOI burn, the solar array broke off. Soon after, aerodynamic torques or aerodynamic heating caused spacecraft loss of attitude control. If the remaining spacecraft was broken apart by aerodynamic pressure or catastrophic rupture of the fuel tanks, the smaller parts would have been captured in the atmosphere. Conversely, if the spacecraft stayed together, aerodynamic drag, coupled with the partial MOI burn, may have provided enough deceleration for capture into Mars orbit. Since there is considerable uncertainty in the calculation, it is difficult to unambiguously predict if the MCO was captured or skipped out of the atmosphere to continue on an interplanetary trajectory. If captured, the body would plunge deep into the atmosphere on the next periapsis pass and break up. In any scenario, only high-density metallic pieces could have survived to reach the surface.

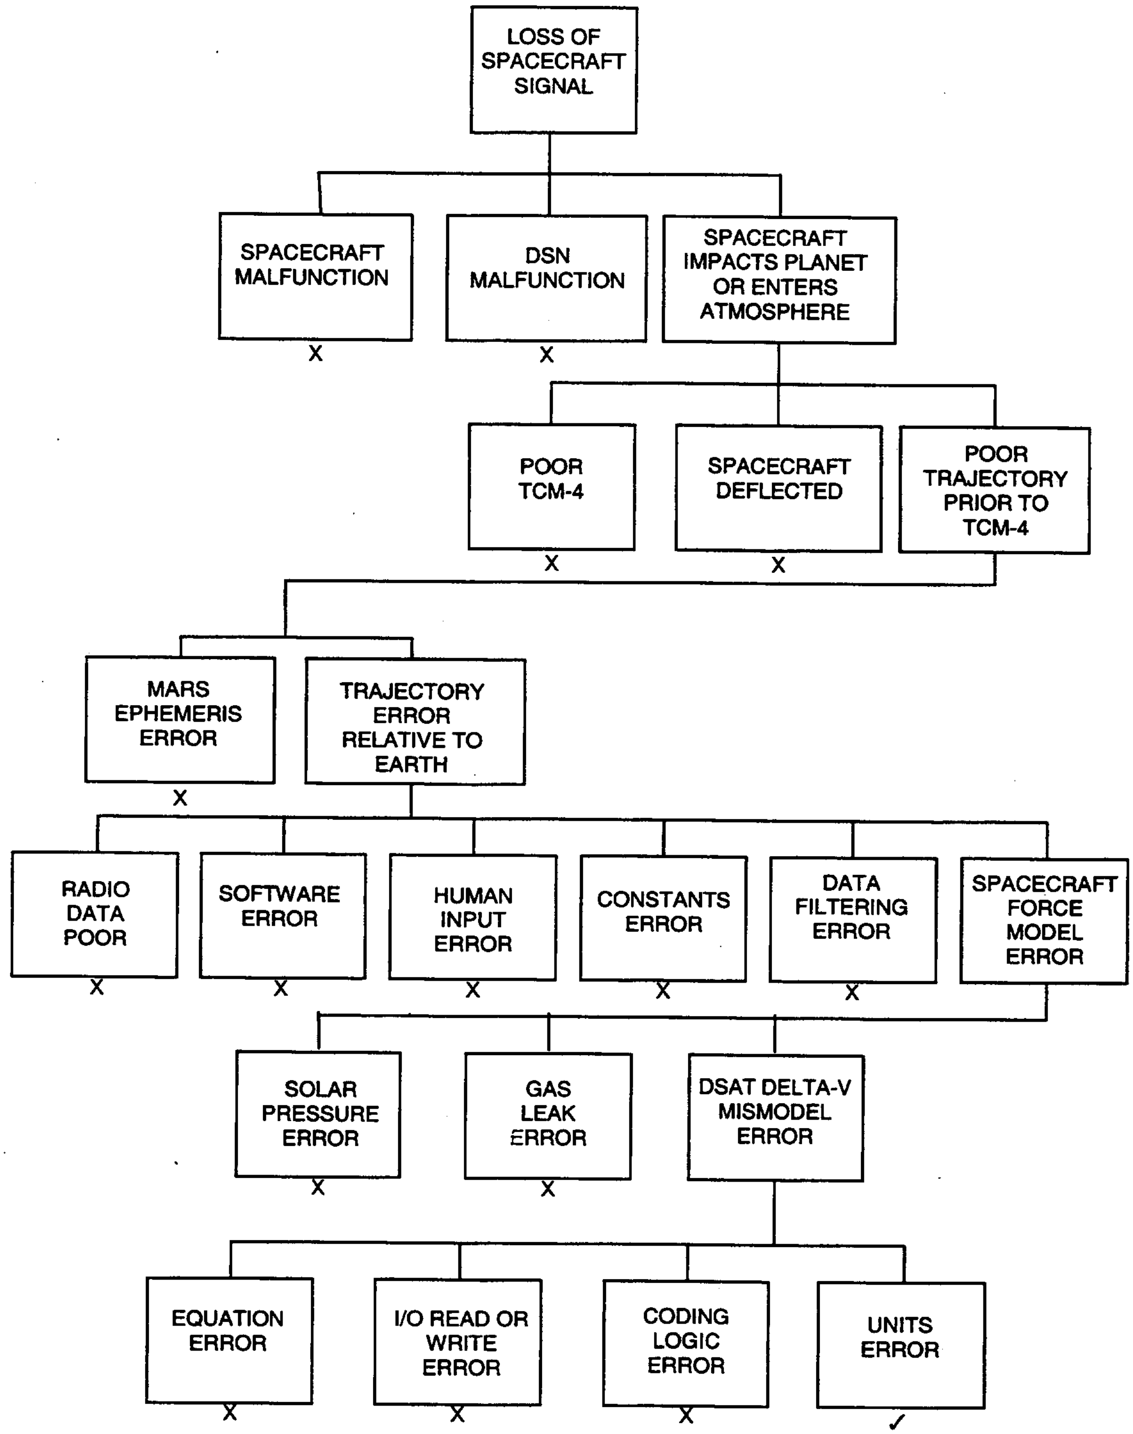

Figure 1: MCO Loss-Of-Signal Truncated Fault Tree Diagram

3 Analysis

3.1.1 Introduction

Section 2 and Appendix 1 detail what happened to the MCO as it prepared to enter the Mars orbital phase of the mission. This section provides an analysis of:

- Why the factors contributing to the loss occurred.
- Why, in spite of established processes and safeguards that have historically mitigated failures, the contributing factors apparently went unrecognized and uncorrected in both the development and operations phases of the mission.

3.1.2 Notes on Figures, Findings and Recommendations

The results of this analysis are presented graphically in Figures 2 through 11, which begin on page 7. Each box in these Figures represents either an event or the absence of an event that contributed to the loss. The organization of the charts is hierarchical. For example, the top box on Figure 2, labeled "Entered Mars Atmosphere", represents the loss of mission statement. The figure shows three causative factors:

- The spacecraft was on the wrong trajectory,
- The fourth trajectory correction maneuver (TCM-4) was wrong,
- The contingency maneuver (TCM-5) was not performed.

These are the reasons why the MCO entered the Mars atmosphere. Each of these three boxes, in turn, has arrows pointing to them that represent contributing factors.

Links are provided between the figures, allowing the reader to move to lower and lower levels through the analysis. At the root levels of the charts are the processes. The "process" boxes have heavy borders to distinguish them from the "event" boxes. The misinterpretation of these established processes, or problems in the MCO project's application of these processes, prevented the MCO team from finding and taking steps to correct against the contributing factors.

The numbers in the boxes refer to the Board's findings, which are included in Section 4. Figures 2 through 11 and the findings constitute the "root why" analysis. The findings are segregated by process or by topic, as appropriate. A set of associated recommendations follows the findings for each process or topic.

3.1.3 The Spacecraft was on the Wrong Trajectory

The MCO was on the wrong trajectory because the MCO navigation process did not correctly account for the small forces acting on the spacecraft.

Two separate but related factors contributed to the error in small forces modeling:

- The omission of the English-to-metric conversion factor in the software program used to generate the AMD files.
- The LMA software program used to generate the AMD file also incorrectly formatted the data on the files, rendering the files unreadable by the navigation software. The format problem was not corrected until mid-April 1999, with the result that the small forces inconsistency was not uncovered until more than four months after launch.

See Appendix 3 for additional information regarding small forces error modeling.

The reasons why the small forces modeling was wrong and remained undetected, either in the pre-launch development phase or in the operations phase, are illustrated in Figures 4, 5, and 6.

Because the small forces modeling was wrong, there were discrepancies between:

- The magnitudes of the small forces calculated from information provided by the Spacecraft Team in the AMD files.

  and

- The estimates of the small forces coming from the navigation analyses.

As is shown in Figure 2, the discrepancies contributed both to the spacecraft being on the wrong trajectory and to the Navigation Team miscalculating targeting uncertainties.

The anomalous results yielded by the navigation process prompted the Navigation Team and the Spacecraft Team to investigate the source of the problem. Contrary to established practice, the investigation was not carried to completion. For example, the telemetry data included an onboard calculation of the delta-V resulting from the small forces; had this data been analyzed, the small forces units error would have been quickly detected.

Without explicit knowledge of the reason for the discrepancy, the MCO navigation analyses should have yielded significantly larger uncertainties for its estimate of the target point. This did not happen. As a result, the wrong targeting uncertainties contributed to the other two causative factors shown in Figure 2. How and why Mission Operations did not properly account for the small forces discrepancy is illustrated in Figure 7 and discussed in further detail in Appendix 4.

3.1.4 TCM-4 Was Wrong

The second causative factor was that TCM-4 was wrong (see Figure 2). TCM-4 was the fourth of the planned trajectory correction maneuvers. TCM-4 was conducted on September 15, 1999, a little more than a week before Mars encounter. TCM-4 was wrong because the MCO navigation analysis yielded the wrong trajectory and because the aim point chosen for TCM-4 was too low, given the true uncertainty in the knowledge of the trajectory. (See Appendix 1 for further discussion.)

3.1.5 TCM-5 Not Executed

The third causative factor was the decision not to execute TCM-5. The mission plan called for the execution of a fifth contingency maneuver (TCM-5) in case of trajectory errors following TCM-4. However, this maneuver was never executed. In spite of anomalous orbit determination results following TCM-4, the mission managers decided not to execute TCM-5, the risk-reduction contingency maneuver. Again, a major factor in this decision was project management's lack of knowledge of the true trajectory uncertainty. The other major factor in deciding not to execute TCM-5 was project management's decision to change the baseline risk-management strategy in the days immediately preceding Mars orbit insertion. Appendix 5 discusses the issues relating to changing the risk management strategy after TCM-4.

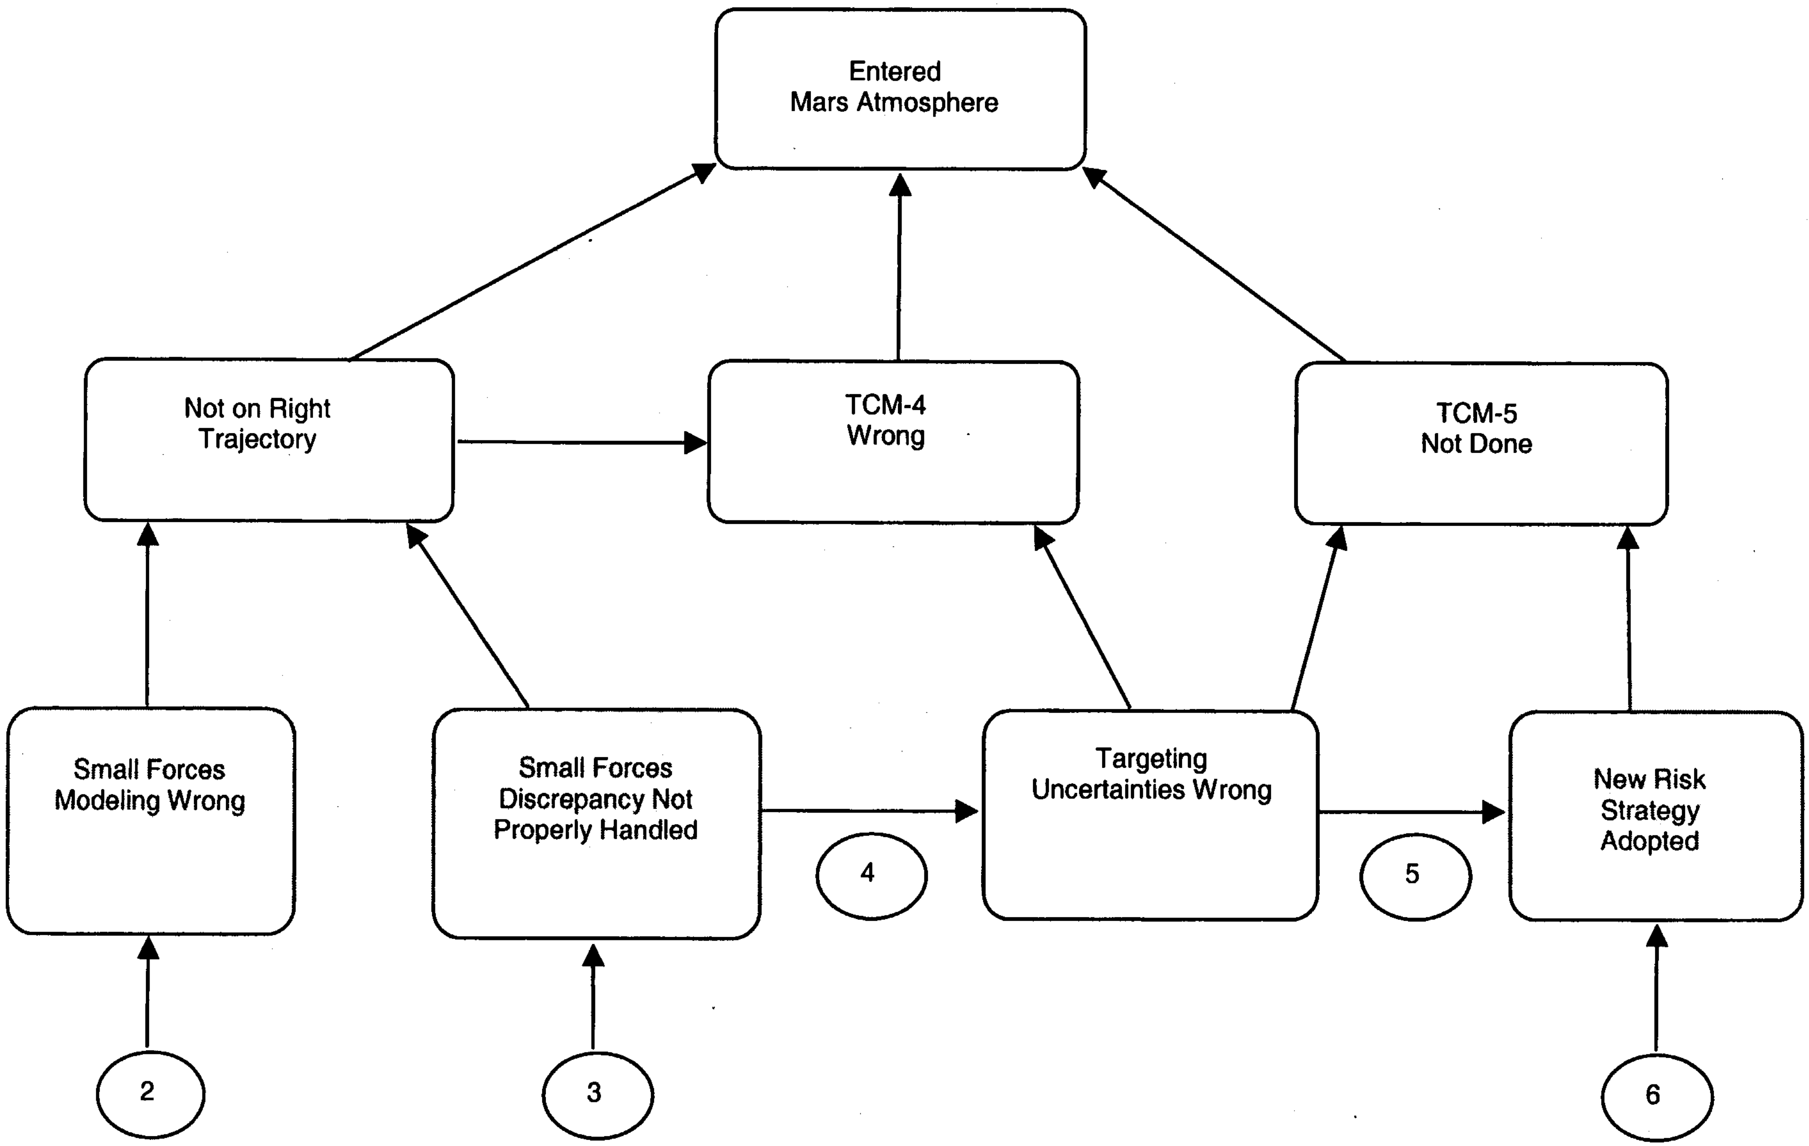

Figure 2: Reasons That The MCO Entered The Atmosphere

18441

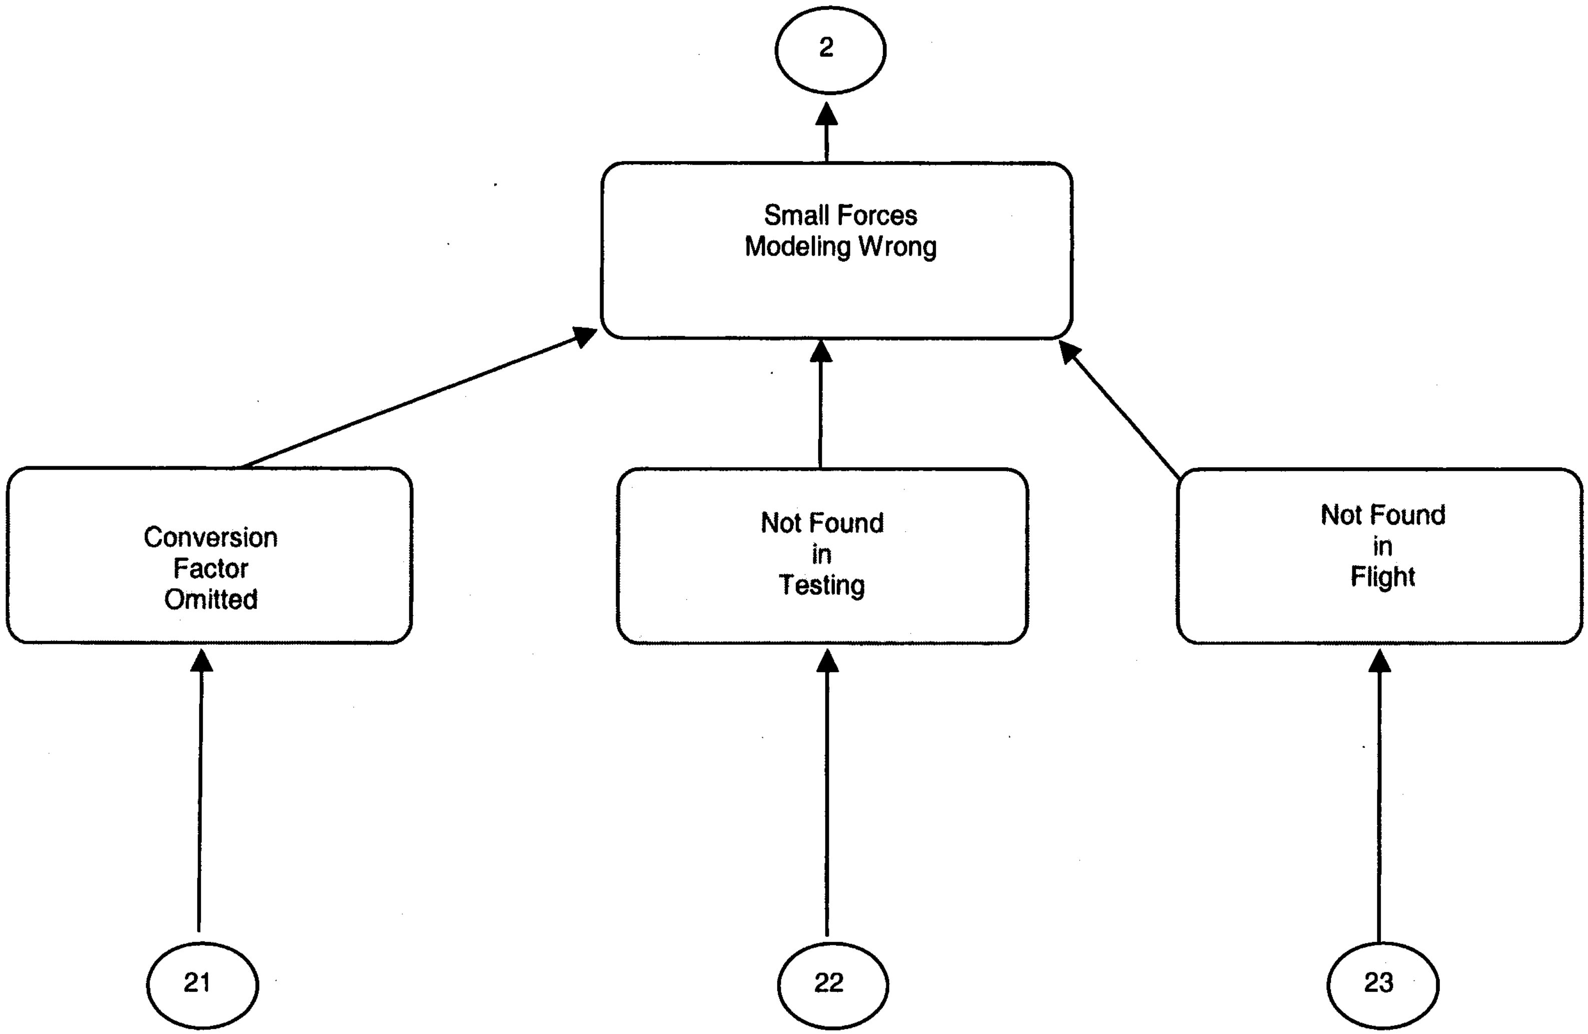

Figure 3: Reasons That The Small Forces Modeling Was Wrong

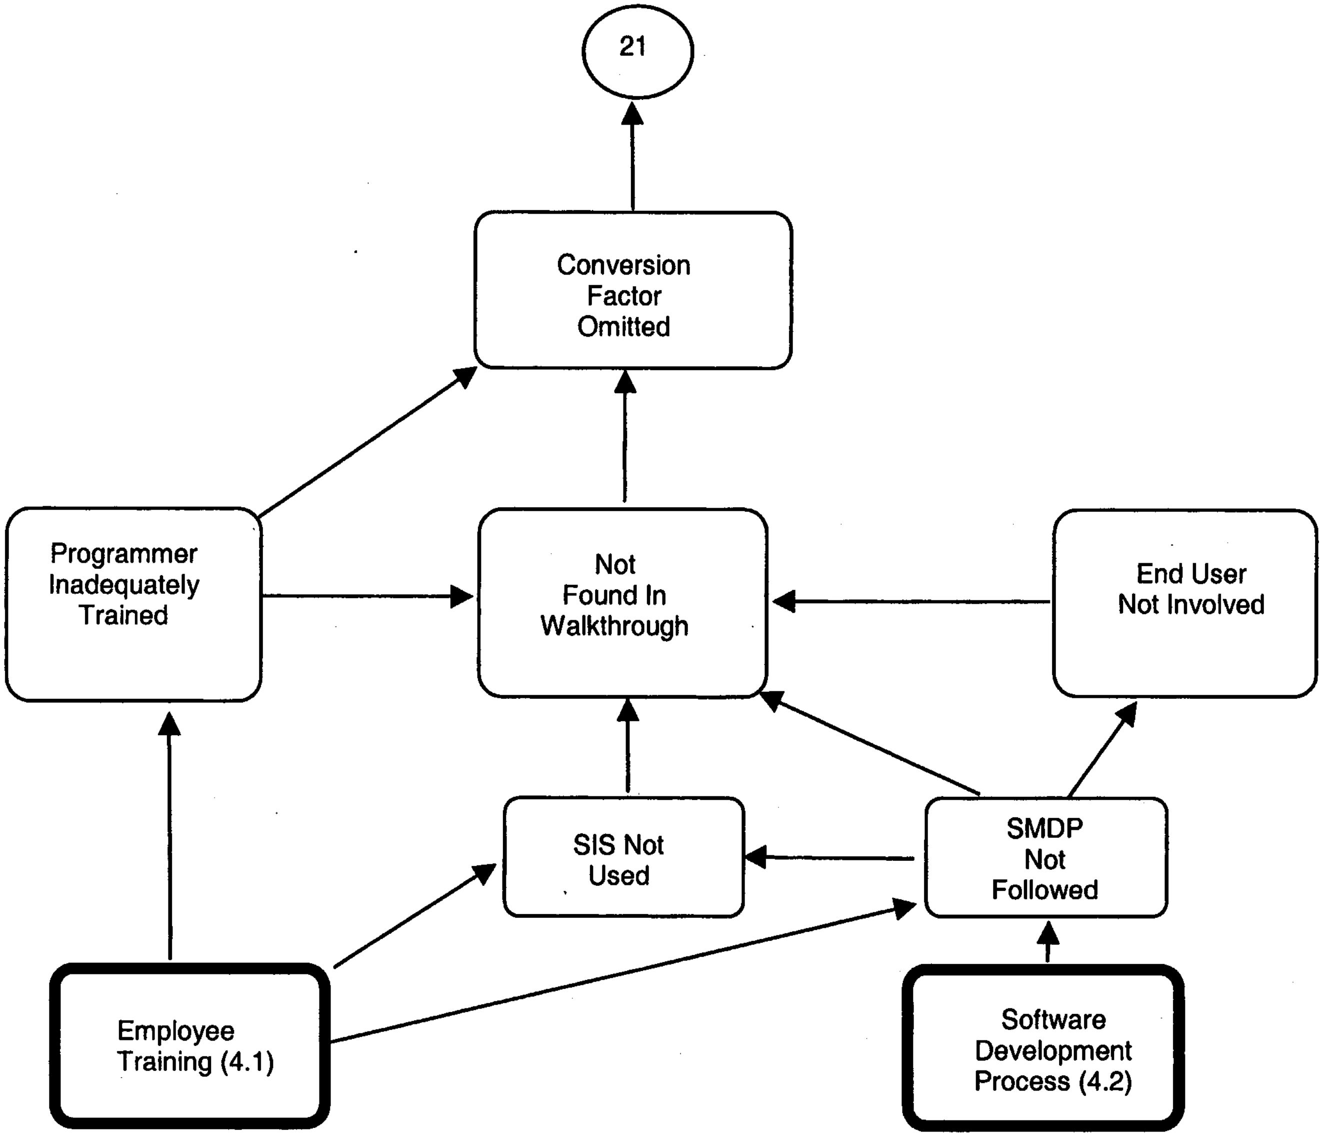

Figure 4: Reasons That The Conversion Factor Was Omitted

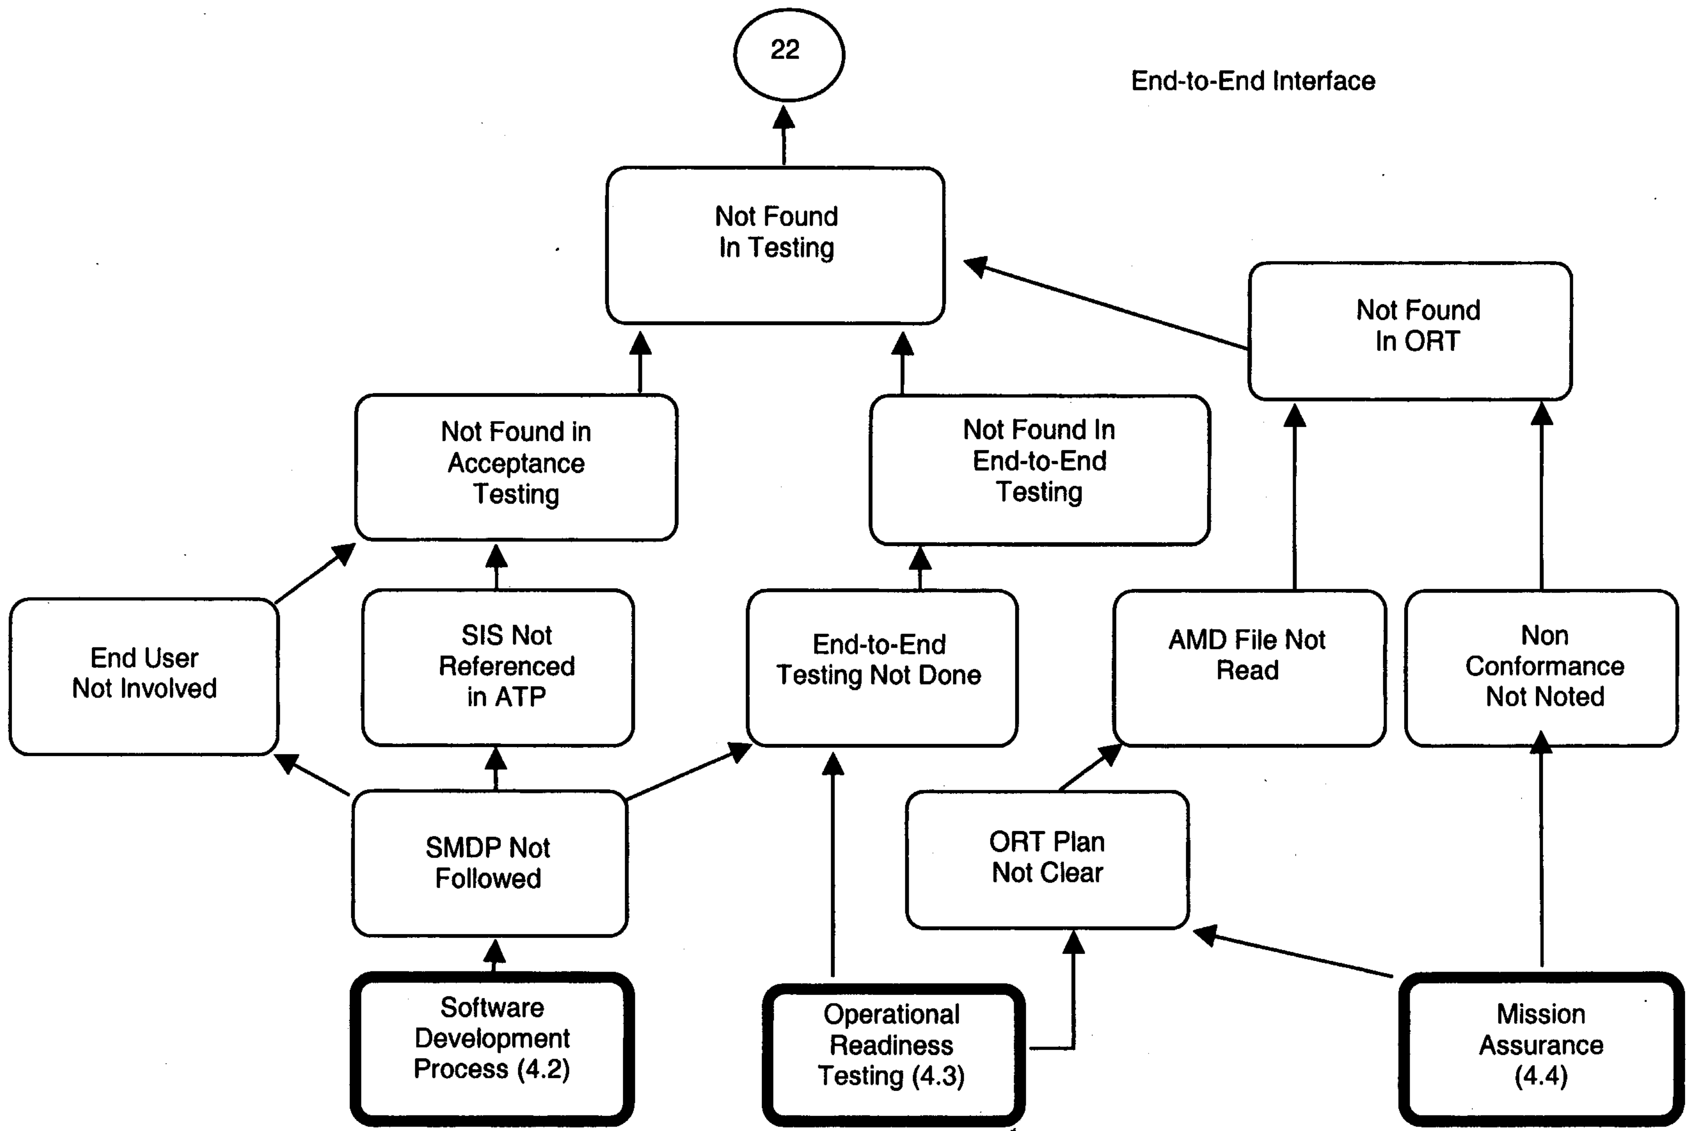

Figure 5: Reasons That The AMD File Problems Were Not Found In Testing

18441

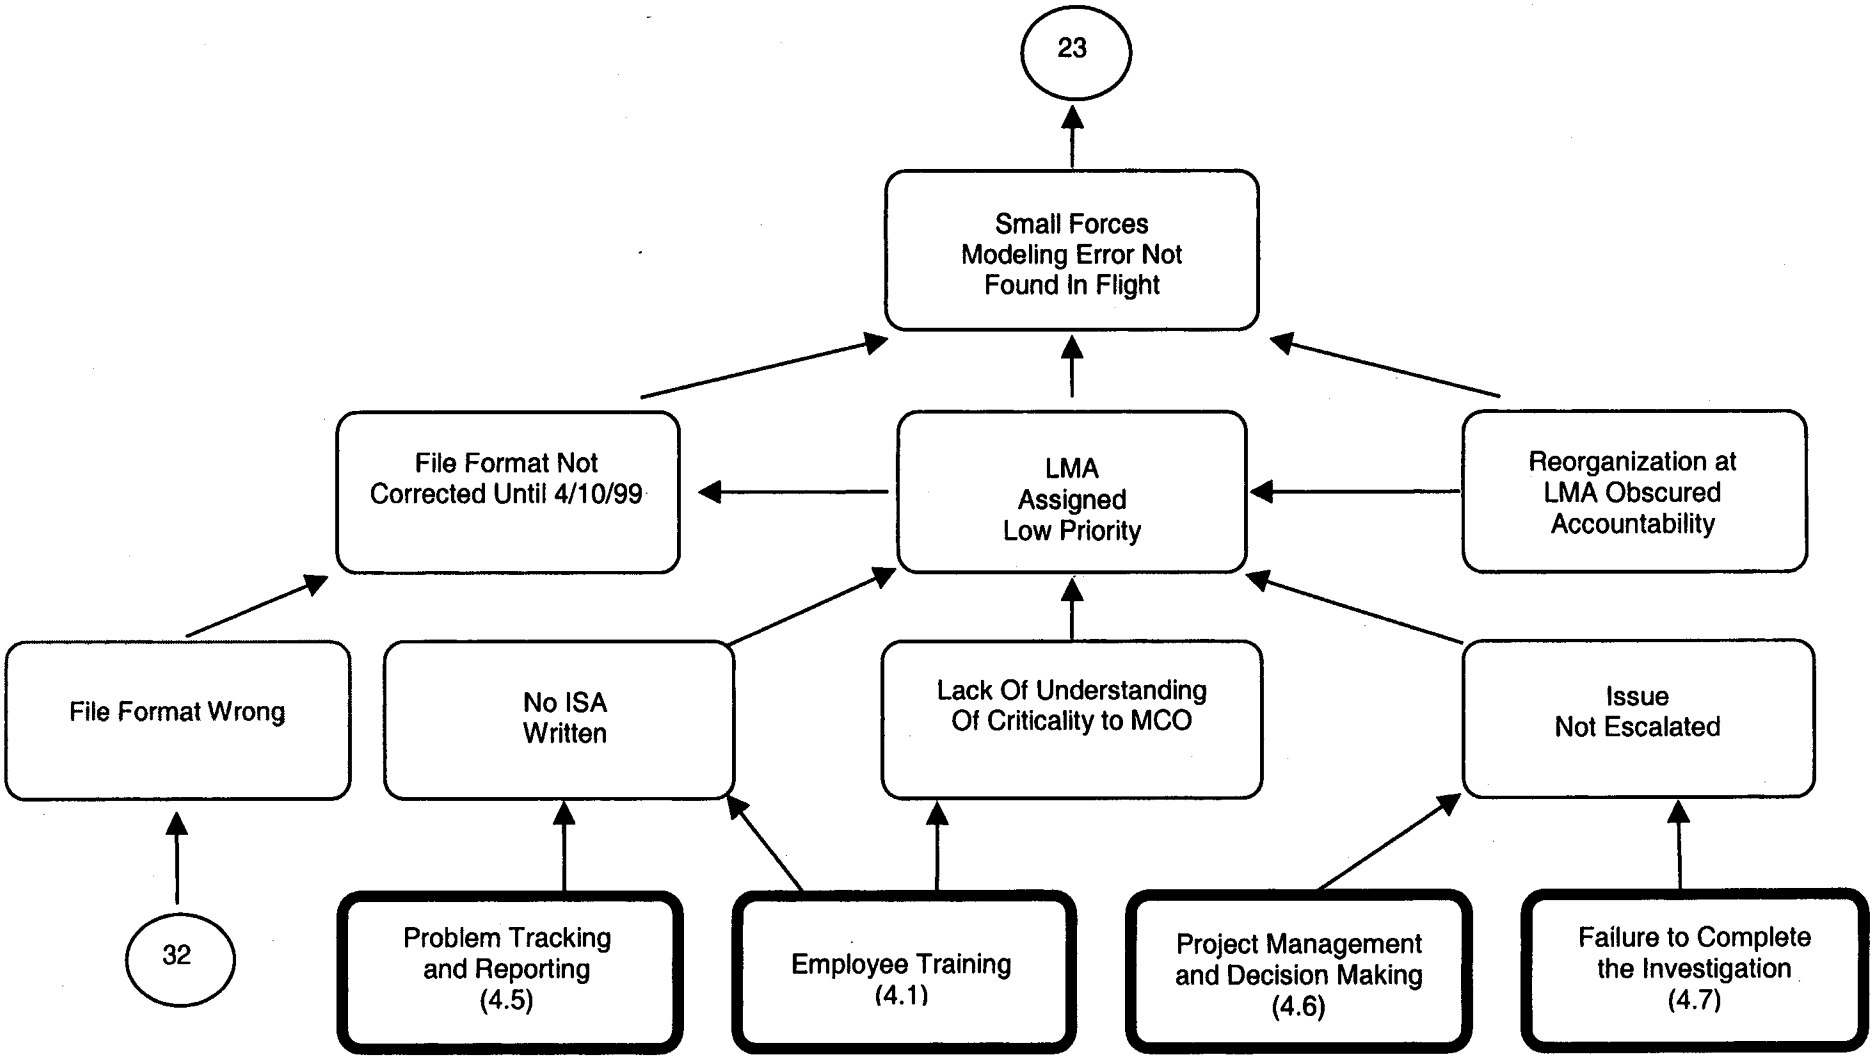

Figure 6: Reasons That The AMD File Unit Problem Was Not Found in Flight

18441

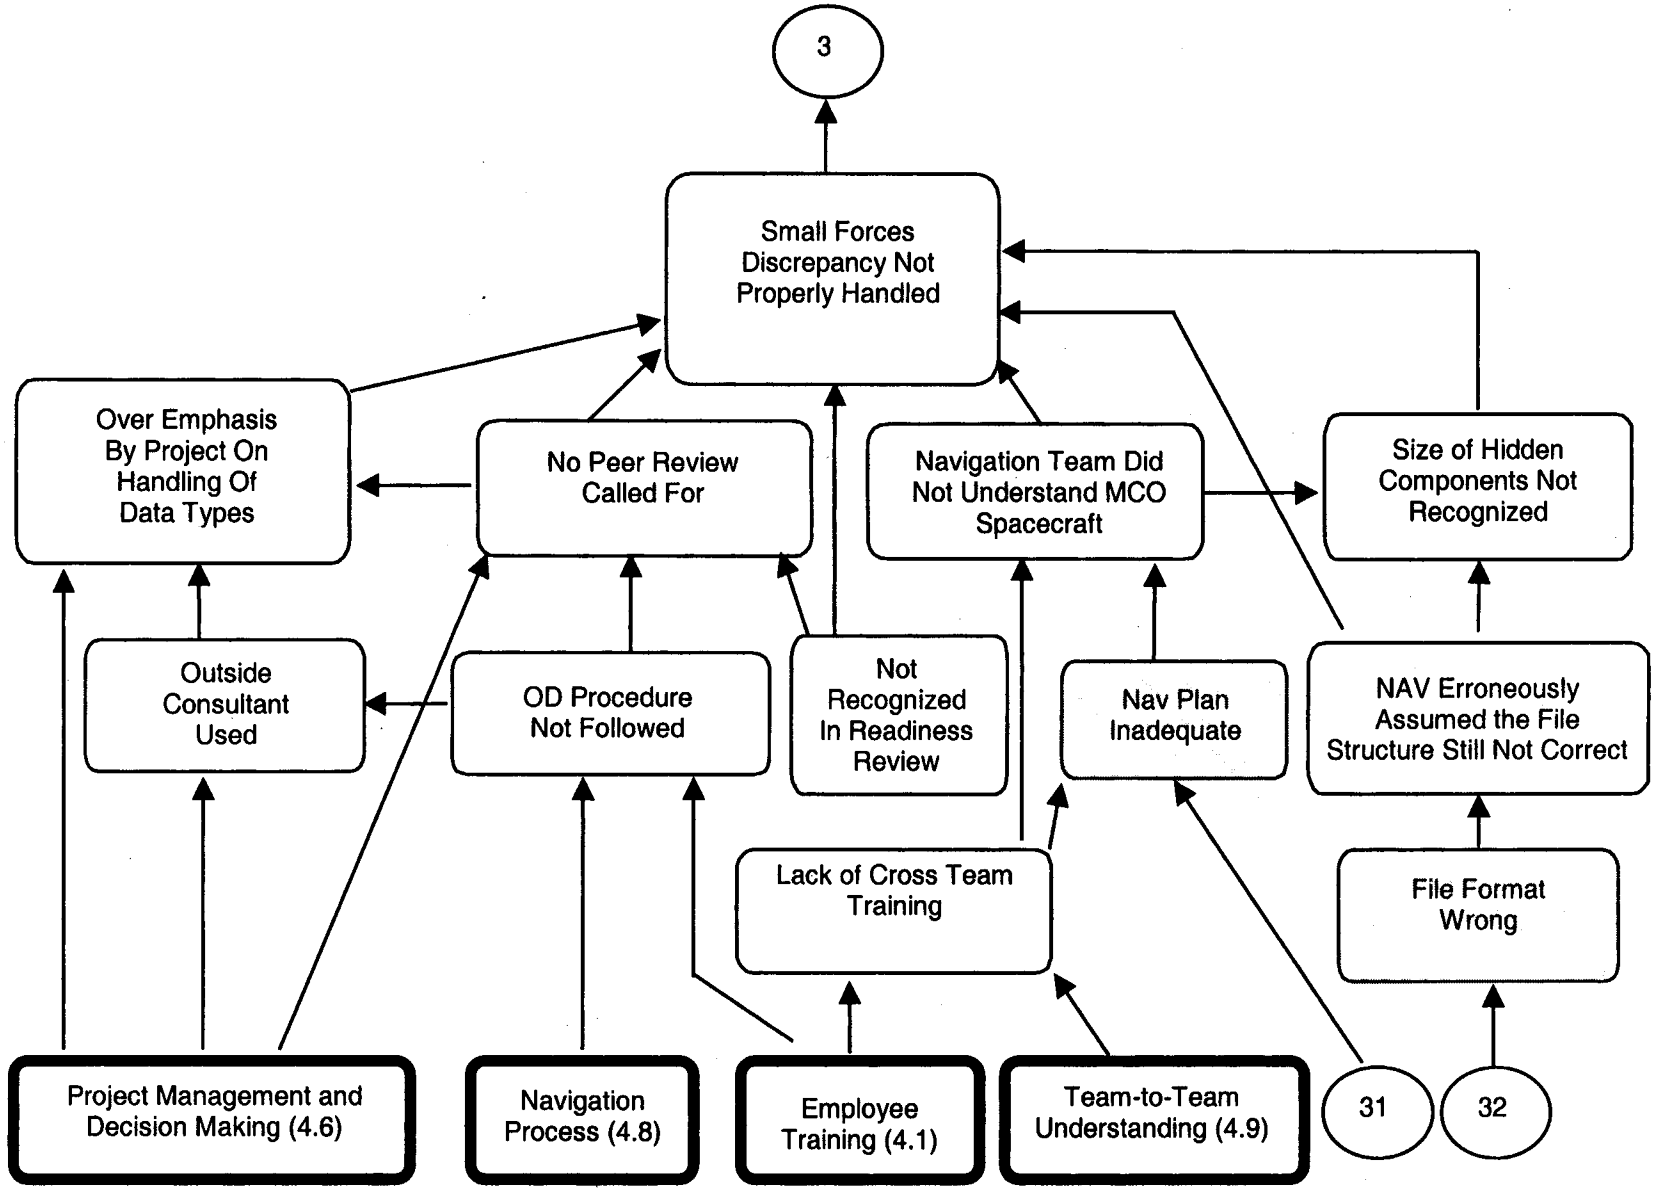

Figure 7: Reasons That The Small Forces Discrepancy Was Not Properly Handled

18441

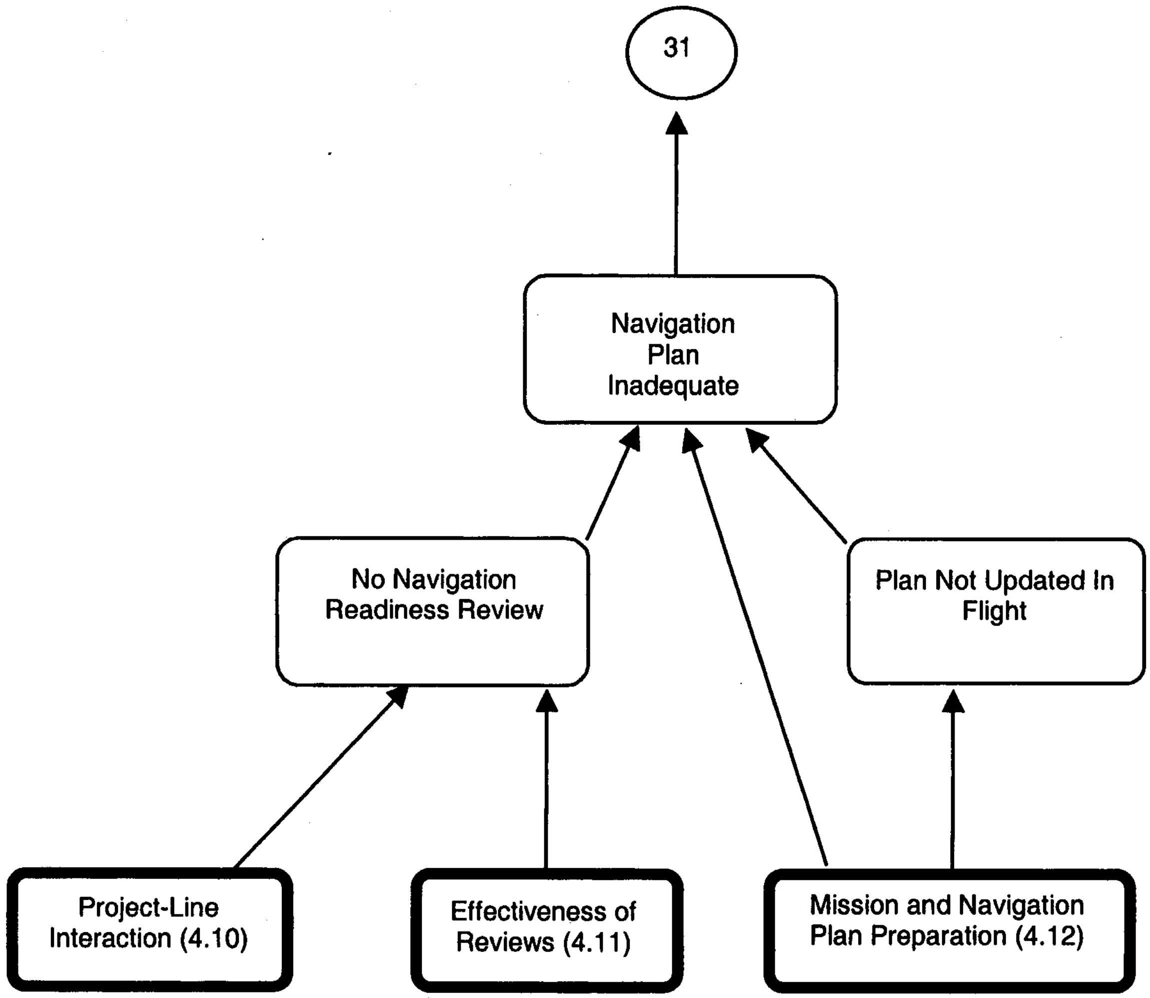

Figure 8: Reasons That The Navigation Plan Was Inadequate

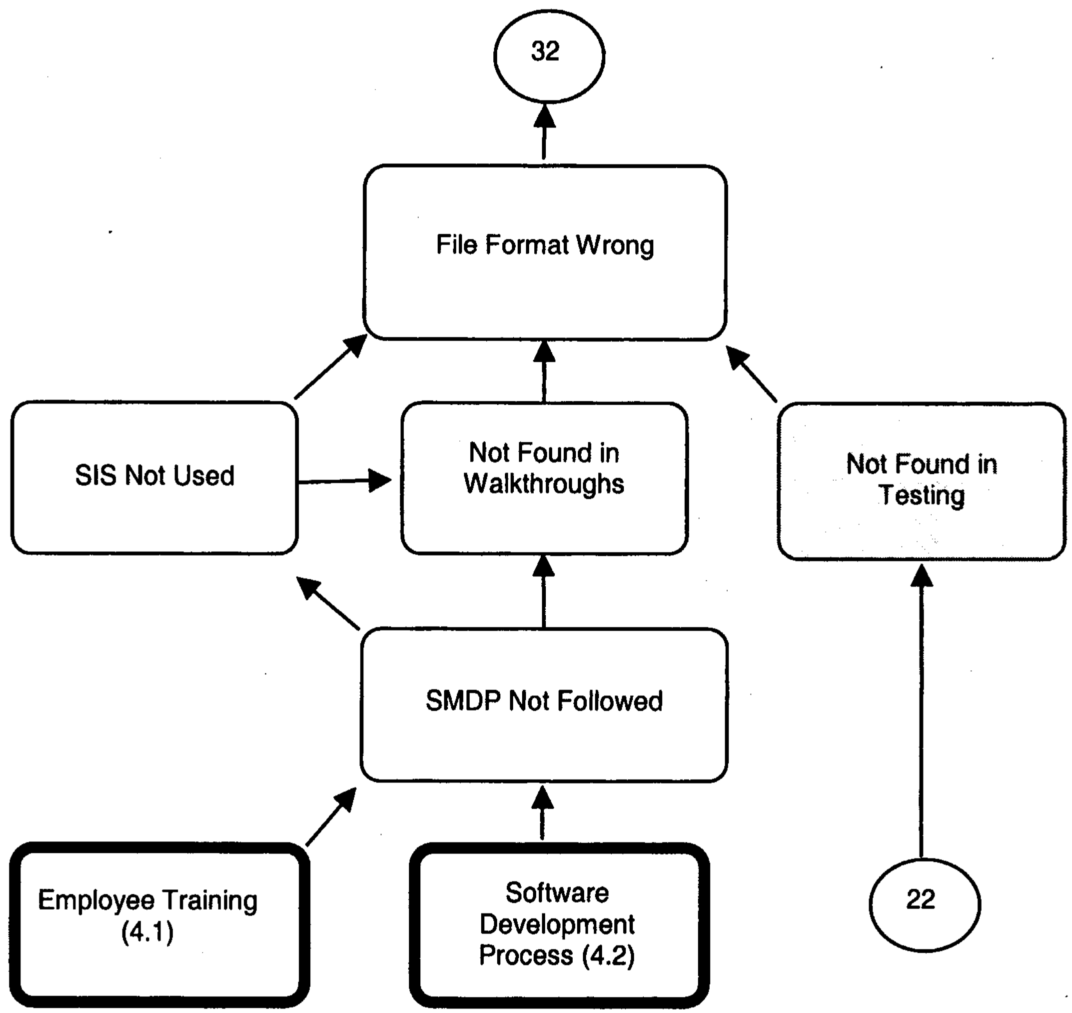

Figure 9: Reasons That The File Format Was Wrong

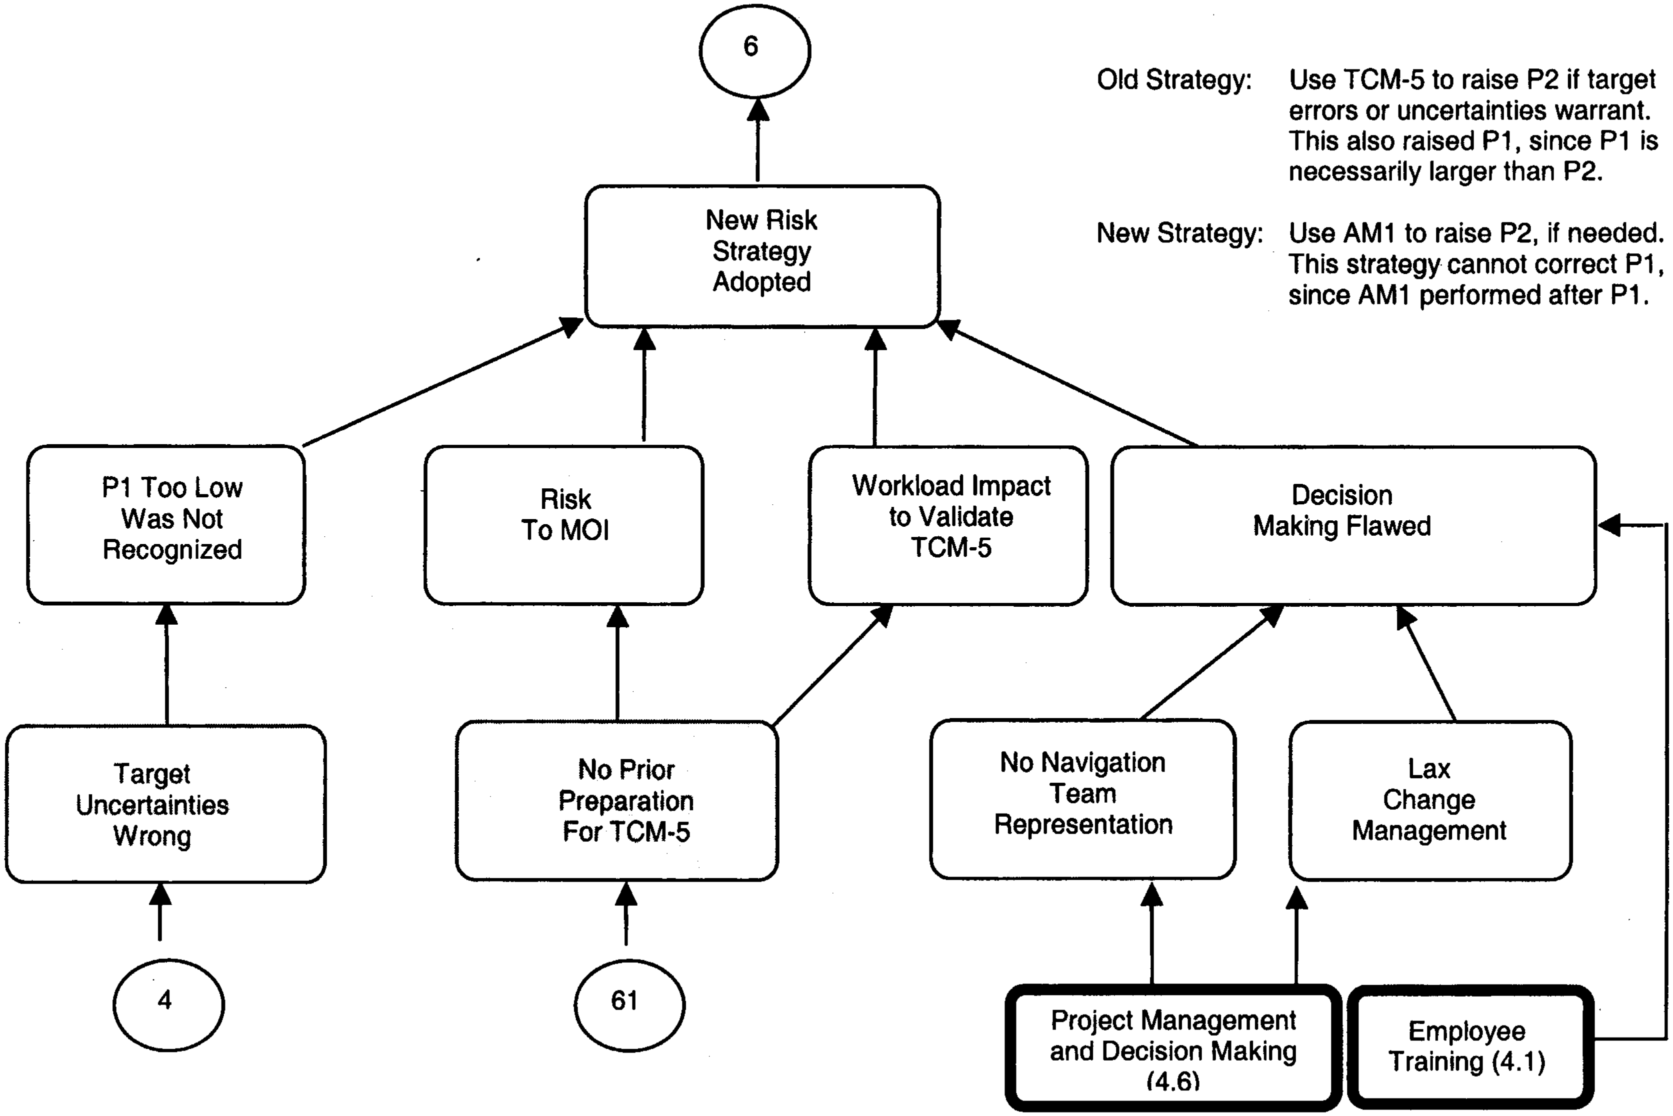

Figure 10: Reasons That The New Risk Strategy Was Adopted

18441

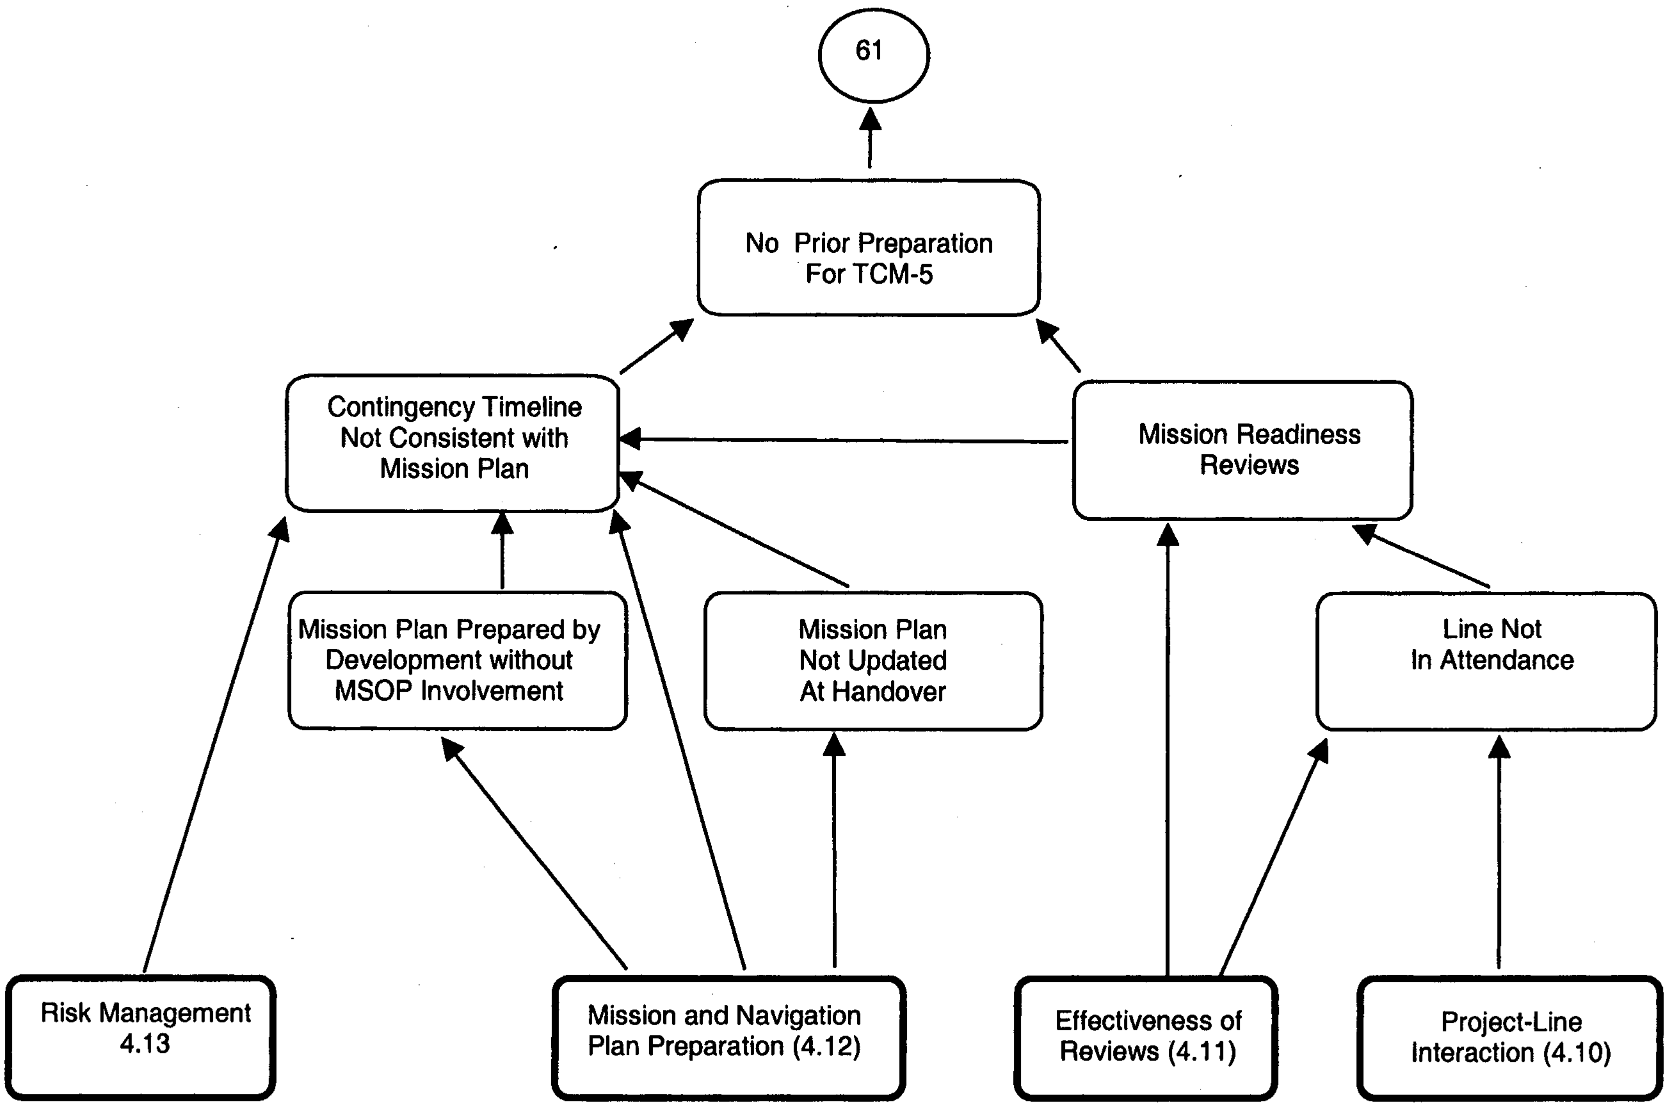

Figure 11: Reasons That There Was No Prior Preparation for TCM-5

18441

4 Findings and Recommendations

4.1 Employee Training

4.1.1 Findings

There were instances where the required training for individuals was not identified, where inadequate training was provided, and where verification of the required understanding was not accomplished. In addition, there was insufficient cross-team training (see section 4.9).

4.1.2 Recommendations

R1) Each employee must be provided with the required training such that they have a working understanding of the procedures needed to perform their assigned tasks and why adherence to the procedures is necessary. **(MPL)**

R2) The Laboratory policies and procedures pertaining to training for line and project management-level positions must be reviewed for applicability to the current environment of many small parallel projects.

R3) Procedures that call for peer reviews must refer to the policy on Reviews (see section 4.9).

4.2 Software Development Process

4.2.1 Findings

The software program called "Sm_forces" had an engineering units error in the calculation of the thruster impulse. Specifically, the AMD file generated by the Sm_forces program provided data to the navigation software in pounds-force seconds, instead of the required Newton-seconds. This error under-represented by a factor of 4.5 the actual size of the delta-V resulting from AMD events. The Sm_forces program also produced AMD files in an incorrect data format, making the files unreadable by the navigation software. MCO personnel did not find these errors in the requirements, design, code, or test walkthroughs.

Contrary to typical practice, there was not a requirement for a navigation (end-user) representative to be present at any of the walkthroughs or the acceptance test.

The Sm_forces software program was misclassified as non-mission critical, which reduced the number of reviews done on the software compared to mission critical software.

LMA did not follow their Software Management and Development Plan (SMDP) in the development of the Sm_forces software program. Specifically, required persons were not in attendance at the walkthroughs and the Software Interface Specification (SIS) was not used in the walkthroughs. Also contrary to the SMDP, there were no minutes of the walkthrough meetings, action items from the review were not published, and the interface with navigation software was not tested. Further, there was inadequate training of software and other project personnel in the software walkthrough process.

The flowdown of requirements from the higher-level MCO document (the SIS) to the LMA software requirements documents was not done. There was no recognition of parameter engineering units in the Sm_forces Software Requirements Specification other than a note to look at the SIS.

4.2.2 Recommendations

R4) The Review Board recognizes that different degrees of formalism of the software development process may be appropriate depending on the criticality and end use of any given piece of software. Mission Critical Software requires the highest degree of formalism. Rigorous criteria must be established for determining the mission criticality of software. As a minimum, Mission Critical Software must include all flight software and all software used in the uplink and navigation processes. **(MPL)**

R5) The end user of mission-critical software or any software that is in the interface between two systems or two major organizations should be required to participate in its requirements, design and acceptance test plan walkthroughs, and to review acceptance test results. This participation should include active involvement in generating the success criteria for the acceptance test, plus approval of the acceptance test plan and acceptance test results. **(MPL)**

R6) For Mission Critical Software, the Board has recommendations in three areas: **(MPL)**

- Software walkthroughs:
  - Update the Software Management and Development Plans at JPL and LMA to ensure that the walkthrough process verifies that the software will provide proper engineering units on all parameters during requirements definition, design, code, and test walkthroughs.
  - Train all software and supporting personnel in the proper conduct of software walkthroughs, including requiring (1) that the minutes of all software walkthroughs are available to project management and to the software line organization manager for review, (2) the tracking of all action items from all walkthroughs until they are closed, and (3) that the status of open actions be reported during Flight Team status meetings.
- Testing (see 4.7 for additional findings and recommendations on end-to-end testing):
  - Conduct performance and end-to-end functional tests on all software program-to-program interfaces prior to launch. Provide the status of that testing and the test results at the launch readiness review.
  - Test all interfaces between the ground system and the spacecraft from end to end (i.e., from command generation through the spacecraft interface to receipt and processing of telemetry).
  - Test all interfaces between the Spacecraft Team and Navigation Team software in an end-to-end sense. Use the System Test Lab (STL) to simulate the spacecraft and its environment as a way to validate expected spacecraft response (for example, delta-V imparted by small forces).

- Software Problem Reporting (see section 4.5 for findings and recommendations regarding ISA reporting):
  - Use the software problem reporting system to document any observed or suspected software problems. Review of the problem report by the software manager should include careful consideration of raising the documentation to a Problem/Failure Report (P/FR) before launch or to an Incident/Surprise/Anomaly (ISA) report after launch.
  - After launch, the ISA must be used whenever the software has a problem in an interface to another flight team organization.

4.3 Operational Readiness Testing

The objective of ORTs is to demonstrate that the operations personnel are trained in the use of the procedures, that the procedures are correct, and that the people using the procedures can achieve the objectives of the test within the required timeline and in the face of unforeseen operational conditions. ORTs usually presume demonstrated functionality of hardware and software interfaces, and are usually run with valid or realistic operational data.

The objective of End-to End testing is to demonstrate interface compatibility among different elements of a system. The demonstration need not be conducted against a timeline and may well involve test data (as opposed to meaningful operational data).

Care must be taken when combining end-to-end test objectives with ORTs to be sure that both sets of objectives can be achieved without compromise to either.

4.3.1 Findings

Contrary to established practice, a specific end-to-end test to validate AMD file compatibility with the Navigation software was not performed. Instead, the intent was to accomplish AMD file validation during the pre-launch Operational Readiness Test (ORT). The AMD file was placed on the File Interface Server in preparation for the ORT, but it was never read or accessed as part of the ORT. Subsequently, there was no ORT non-conformance report to indicate that a planned objective of the ORT had not been accomplished. Relative to other flight projects, the ORT was not conducted as rigorously as it should have been in assuring that all interfaces were tested with real data.

4.3.2 Recommendations

R7) ORT test scripts must be reviewed for completion with respect to incorporated end-to-end test objectives. **(MPL)**

R8) Mission assurance oversight must include all objectives of ORTs. Non - conformances must be recorded and tracked for all ORT objectives. **(MPL)**

4.4 Mission Assurance

4.4.1 Findings

Current Laboratory procedures give managers of projects in the operations phase the discretion to choose whether or not to have a project-level mission assurance function. Although there was a mission assurance function established at LMA, the Mars Surveyor Operations Project (MSOP) had no mission assurance function for the work performed at JPL. A number of mission assurance lapses were noted. For example:

- Contrary to established practice, the spacecraft was launched without establishing that the ground-data system navigation software and the AMD files were compatible in format or correct in content. Also, counter to typical practice, the Operational Readiness Testing was not comprehensive in the definition of objectives and criteria for success.
- Problem reporting of the incompatibility of the AMD file format and delta-V inconsistencies were only informally documented in e-mails or memos.

4.4.2 Recommendations

R9) The Laboratory mission assurance policies and procedures must be modified so that all projects in operations must have a Mission Assurance Manager. **(MPL)**

R10) The mission assurance function for all projects in operations must include broad, independent responsibility to verify the adequacy, completeness, and compliance with both design and operational requirements. **(MPL)**

R11) For each project, assign an experienced person charged with overall mission success oversight during mission operations. This person would ideally have participated in the development project. **(MPL)**

4.5 Problem Reporting, Tracking, and Closeout

4.5.1 Findings

The JPL Standards Document entitled Problem/Failure Reporting System Guideline and Procedures provides explicit guidance on the use of P/FRs in the development phase. There is not a similar document covering ISA reports, Anomaly Reports, Discrepancy Reports, and Failure Reports used during operations for addressing post launch issues.

The Review Board believes that the lack of such a document contributed to the following:

- The Failure Report on AMD file format problem was not generated until on April 6, 1999. The problem was observed December 12, 1998. No ISA was written.
- A discrepancy between the delta-Vs expected by the Navigation Team and those produced by the AMD file was only identified in an e-mail on April 26, 1999. The Navigation Team should have written an ISA.

Although not a contributory factor in this instance, it was found that the Spacecraft Team's practice was to have the Chief Engineer write all of the ISAs. This practice has the desirable effect of pre-filtering and perhaps precluding the generation of an unnecessary ISA. On the other hand, introducing an additional step and person into the ISA process may result in a necessary ISA not being written.

4.5.2 Recommendations

R12) Projects must rigidly enforce the requirement to write ISAs. All project members should be empowered to write an ISA without management or other encumbrances. **(MPL)**

R13) The owner of the JPL Problem Reporting and Tracking and the LMA equivalent must update the JPL Standards Document Problem/Failure System Reporting Procedures, D-8091 to include post-launch problem reporting and tracking. This document needs to address critical risk identification and formal reporting requirements.

4.6 Project Management and Decision Making

4.6.1 Findings

Contrary to typical practice, the JPL MCO project development team did not ensure that the Navigation Team became familiar with those spacecraft characteristics that could have influenced navigation performance. For example:

- Members of the MSOP Navigation Team were not invited to major development phase reviews (e.g., Preliminary Design Review and Critical Design Review).
- The Navigation Plan was not updated to accurately reflect the desaturation frequency or delta-V magnitudes (see section 4.12).

As a consequence of an inadequate navigation performance plan (see section 4.12), the Navigation Team was surprised after launch. However, no ISA was written, as required by the MSOP Configuration Management Plan, and the project did not undertake to re-do the navigation analysis, thereby leading to continued underestimation of trajectory uncertainties.

The MCO project did not pursue either the small forces discrepancy issue or the orbit determination issues to a point that would have forced resolution. The relationship between the small forces and orbit determination problem was recognized. MCO navigation and project management believed that the risk management strategy (raising the aim point as necessary to keep P2 above 150 km 3 sigma (see section 4.13)) would compensate for orbit determination uncertainties.

Project management did not seek an outside peer review of the orbit determination problem (see section 4.9). Instead, a single navigation consultant was commissioned. The consultant focused on data quality issues.

A hastily organized meeting of a small subset of the MSOP Standing Review Board convened on September 18, 1999, to discuss the need for TCM-5. Contrary to the typical practices of other flight projects, the Navigation Team was not represented at this meeting. The utilization of a small subset of the Operational Readiness Review Board to make recommendations on operational decisions without all viewpoints is inappropriate.

MCO project management did not communicate to the Navigation Team a major change in the TCM-5 go/no-go decision strategy. Specifically, the Go/No-Go meeting scheduled for Sunday, September 19, 1999, was superceded by the provisional decision made Saturday morning described above. The project management, in consultation with the Spacecraft Team and the attending members of the Review Board, made a provisional decision not to proceed with TCM-5 provided the navigation results were substantially the same as those of the previous day. The Navigation Team was not informed of this change in decision-making strategy.

4.6.2 Recommendations

R14) The project management training process should be strengthened to emphasize:

- Decision making and evaluating the risk of options.
- The importance of identifying mission critical decisions that could be potentially irreversible or catastrophic.
- The need for all project-level decisions to be documented, communicated to Project members in a timely fashion, and consistent with project Configuration Management Plan requirements. **(MPL)**
- That project-level decision affecting requirements, schedule, resources, and risk should be made with full representation by all project elements with expertise relevant to the decision issue. **(MPL)**

4.7 Failure to Complete the Investigation

4.7.1 Findings

The Navigation team identified errors in the small forces impulse data. During May and June there were several informal e-mail messages back and forth between the Navigation Team and the Spacecraft Team in an attempt to resolve the errors. However, the investigation was never completed because it was not considered mission critical and many "higher priority tasks" were determined to exist. The MCO project did not escalate the issue with LMA.

4.7.2 Recommendations

R15) Requests for actions from one team to another must be documented using a disciplined disposition, tracking, and resolution process. **(MPL)**

4.8 Navigation Process

4.8.1 Findings

The MCO navigation process did not correctly identify the actual trajectory. Orbit determination solutions exhibited a large spread and were changing in a manner that strongly suggested modeling errors in the process. The source of modeling errors consistent with the solution spread was not identified.

The MCO navigation process did not correctly bound the uncertainties in the trajectory. As a consequence, the MCO mission managers did not understand the magnitude of the uncertainty in their knowledge of the trajectory, which, subsequently, led to incorrect decisions regarding TCM-4 and TCM-5.

In investigating the trajectory error, the MCO Navigation Team did not comply with the Orbit Determination Operational Procedure, which requires that a peer review be conducted to help select the best solution for subsequent action. The Laboratory policy on Reviews specifically states that, "Peer reviewers shall not be currently working in the project/major task element under review." The reference in the Orbit Determination Operational Procedure to peer reviews is ambiguous in that it does not refer to the Reviews policy. Absent this reference, and absent training to the contrary, the Navigation Team assumed the Webster definition of "peer" and reviewed themselves.

No one person had end-to-end responsibility for the navigation process, which includes the modeling and generation of the a priori inputs to navigation analysis. This, in a very direct way, contributed to not having a demonstration of interface compatibility, either through an end-to-end test or through an operational readiness test. (See section 4.3.)

4.8.2 Recommendations

R16) Navigation Team personnel must be coached by their line management in the use and understanding of the existing Orbit Determination Operational Procedure (NAV)[1], with specific regard to the use of peer review. **(MPL)**

R17) The process owner for the Orbit Determination Operational Procedure must provide further guidance in the procedure regarding the strategies for additional computer runs that should be carried out prior to initiating a peer review and must clarify the procedure. Specifically, the procedure should be changed to make it unambiguously clear that peer reviews must comply with the Laboratory policy on Reviews and that peer reviewers must be knowledgeable experts external to the project.

R18) Navigation must be defined as an end-to-end process. The process describing navigation and the procedures controlling its operation must include all of the related activities essential to producing navigation mission success. The navigation process owner must be accountable for the quality of all of the inputs to the process, including all data and file interfaces, and for the tests needed to demonstrate that the navigation outputs can be delivered to meet the needs and expectations of the project.

4.9 Team-to-Team Communication and Understanding

4.9.1 Findings

Some members of the LMA Spacecraft Team were largely unaware that the estimated delta-Vs derived from the small forces impulse data are a direct input to the orbit determination process, nor did they understand the sensitivity of the spacecraft trajectory to very small (millimeters per second) velocity errors.

The Spacecraft Development Team did not have an appreciation of how several spacecraft design parameters relate to the interplanetary navigation issues (see section 4.12). The lack of knowledge in the above areas influenced the focus and emphasis taken at LMA in its process to deliver accurate small forces impulse data.

The MSOP Navigation Team members were largely unfamiliar with many of the relevant properties of the spacecraft and the LMA ground data system processing. Lacking this information, the Navigation Team incorrectly assumed the delta-V component uncertainties were small, spherically distributed, and uncorrelated, as represented in the navigation plan (see section 4.12).

The anatomy of the MCO mishap showed many instances where simple communications between people were not followed to their conclusions. Key channels for these communications involved various combinations of the Navigation Team, the Spacecraft Team, project management, and line management. For example, the problem during flight could have been solved with:

- Better Navigation Team-to-Spacecraft Team communications to diagnose the Y-axis force level.
- Better project management-to-line organization communications to gather the needed help.

[1] (http://elias.jpl.nasa.gov:8080/cgi/doc-gw/DocID/29694/?KD=dmie)

- Better Navigation Team-to-project management communications to stress the level of concern.
- Better project management-to-Spacecraft Team communications to get action on the Navigation Team-Spacecraft Team issues being raised.

4.9.2 Recommendations

R19) Cross-knowledge between flight team elements must be emphasized in current and future flights. For example, a designated member or members of the Spacecraft and Navigation Teams should be assigned to become knowledgeable about the other team's issues, needs, and technical methods, and to serve as a technical point of contact. **(MPL)**

R20) Cross-team knowledge must be emphasized through appropriate team-to-team communications, such as system and subsystem orientation seminars and mission operations test and training sessions. **(MPL)**

4.10 Project–Line Interaction

The present era of many small projects leads to a need for small, sharp, fast moving project teams. There is not much margin for error and there are not many spare resources within the project teams to deal with emerging problems. The demands on these lean project teams can lead to on-the-spot prioritizing of tasks (see section 4.7.1) and internal problem solving that increases the risk of missing something important.

4.10.1 Findings

There was insufficient interaction between the MCO project and the line organization. A timely involvement of experienced navigation experts would have revealed the small forces inconsistency or, failing that, should have led to an appropriate characterization of the targeting uncertainty. Contributory factors include a misunderstanding or misapplication of the principles guiding line management involvement in projects (refer to JPL D-12547, Soft Projectization) and a misperception by the MSOP management that line organization involvement would be intrusive and disruptive.

4.10.2 Recommendations

R21) The principles stated in JPL D-12547 must be adopted by policy as applicable to all project and line elements. See Appendix 6 for an excerpt from this document relevant to these findings. This document must be included in the training of all employees and consideration should be given to a new title for JPL D-12547, such as Principles for Project-Line Relationships.

R22) Technical resources external to the projects must be conveniently available to the projects for problem solving and peer review. The Laboratory must cultivate an environment that encourages the use of these resources.

4.11 Effectiveness of Reviews

4.11.1 Findings

There were several instances during both the development phase and the operations phase when the appropriate MCO representatives were not present at reviews. Even when the project review process was followed, the board did not contain the appropriate representatives to identify important issues. For example, the Navigation line organization was not represented at the MOI and Aerobraking Readiness Review during which the TCM-3 error was presented. In addition, contrary to established practice, requests for action were not generated to address issues identified by the board. During the operations phase, neither the project nor the line organization utilized the peer review process as rigorously as intended by typical JPL practice.

4.11.2 Recommendations

R23) Appropriate representatives of the JPL line organizations must be assigned for required attendance for each review type. For example, the Navigation line organization should be assigned to attend all launch readiness and mission orbit insertion reviews. **(MPL)**

R24) Navigation personnel must attend project spacecraft development design reviews.

R25) Review boards must assess the project-proposed review agendas as stated in the Laboratory policy on Reviews. The board chairman must review with the board the criteria for reviews. The review process owner should review the guidelines and procedures with review board chairman. **(MPL)**

R26) Review board members must be advised by the board chairmen to pursue reasons for all anomalies revealed during reviews. **(MPL)**

R27) Laboratory policy and Guidelines for Reviews must be reviewed and recast as necessary to enforce a mandatory procedure and process for peer reviews and to provide improved guidance regarding the scope and purpose of Operations Readiness Reviews (ORRs).

R28) The Laboratory must restate and enforce a mandatory procedure and process for peer reviews. Items that should be included in the peer review process include:

- Information emphasizing that the cardinal objectives of having peer reviews is to infuse the collective expertise of the line organizations into the projects and to promote a collective sense of responsibility for success.
- A better definition of what a peer review is, primarily to distinguish from other more formal types of reviews.
- A description regarding how to make peer reviews “quick turnaround” and non-intrusive to the implementation of a project.
- An established Laboratory practice to fund peer reviews.
- The prerogative for line organizations to call for peer reviews when appropriate.

- A statement that peer reviewers must be:
  - Technical experts.
  - Drawn from a talent pool selected by the line organization.
  - External to the project under review.

4.12 Mission and Navigation Plan Preparation

4.12.1 Findings

Contrary to established practice, the Mars '98 Project did not update the navigation accuracy analysis prior to launch to reflect the spacecraft design characteristics. The analysis was in error because it was based on one desaturation maneuver per week, rather that the expected 2 desaturation maneuvers per day. The analysis included no prediction from the spacecraft designers regarding what the desaturation delta-V direction and magnitude would be during the mission. However, the analysis did acknowledge that a more accurate modeling of small forces would be required to meet the 150-km minimum P2 altitude objective.

Contrary to typical practice, the MSOP Navigation Team did not participate in the pre-launch navigation accuracy studies and did not re-do the navigation accuracy analysis, either in preparing to take over the mission or after the desaturation frequency was revealed to be two times per day, instead of once per week.

Spacecraft design and trade studies were undertaken during development without the benefit of input, review, and assessment from the Navigation and Flight Mechanics Section. Examples include, the choice of unbalanced solar array and uncoupled desaturation thrusters, the choice of thrust direction to perform desaturation maneuvers, and the introduction of the "barbecue" mode and the subsequent dropping of this mode.

4.12.2 Recommendations

R29) The navigation and mission plans created by the development organization must ensure operations involvement during approval and handover and must be up to date and complete at the time of handover to the operations organization.

R30) Operational factors must be taken into consideration during spacecraft design and development trade offs. For example, in the case of navigation, this would include delta-Vs from thrusters, torques from unbalanced solar pressure, calculation of spacecraft effective area, and mission operational procedures and timelines.

R31) In-flight updates to these plans should be subject to both project and peer review. **(MPL)**

4.13 Risk Management

4.13.1 Findings

The mission plan called for TCM-5 two days before Mars Orbit Insertion (MOI), if needed. However, TCM-5 had not been included in the mission operations timeline. Specifically, TCM-5 had not been tested for compatibility with a running MOI sequence during the course of implementing and preparing for MOI. Neither was TCM-5 part of the Operational Readiness Testing. As a consequence, when the time came to execute TCM-5, it became a major risk decision for the project to make at a critical point in time. The project waived the requirement of 150 km (3 sigma) altitude at P2 without following its established change management procedures.

4.13.2 Recommendations

R32) Current and future projects must review their operational scenarios and mission timelines for consistency with their Mission Plans and to determine that the necessary planning is in place to support their risk management strategies. **(MPL)**

R33) When circumstances or events dictate departing from the established plan the management team must follow the appropriate change management procedures, including performing necessary reviews, and take the steps necessary to keep all members of the project aware of the changes.

Appendix 1. The Peer Review Team Investigation

The MCO mission was designed to:

- Be a scientific observatory, orbiting Mars in a polar orbit and using a set of scientific instruments to probe the weather and climate of Mars.
- Serve as the primary telecommunications relay link for the Mars Polar Lander (MPL) for three months after the MPL landing on December 3, 1999 (Mars Global Surveyor (MGS) was intended to serve as the back-up relay link).

The spacecraft was designed to fire its rocket engine near closest approach to Mars to slow the spacecraft so that it could enter orbit. This maneuver was called the Mars Orbit Insertion (MOI) maneuver. MOI was scheduled for September 23, 1999. The spacecraft was targeted for a first closest approach (periapsis) altitude of 226 km using a Trajectory Correction Maneuver (TCM) at MOI minus 8 days (September 15, 1999). The last orbit determination estimate two days prior to MOI was 150 km.

The design of the MOI maneuver was such that 5 minutes after the start of the 16-minute MOI burn the spacecraft would pass behind Mars as seen from the Earth. Twenty-one minutes later, the spacecraft would then reappear to Earth view, having achieved Mars orbit.

On September 23, 1999, the spacecraft began the MOI maneuver by turning to the correct attitude for the burn. Telemetry indicated that all elements had performed as expected. At the expected time, the MOI maneuver started. Tracking data indicated that the acceleration to the spacecraft was as expected. However, the spacecraft passed behind Mars about 49 seconds earlier than anticipated and never reappeared.

Two hours before the start of the MOI, an orbit was computed based on tracking data taken up to 12 hours before MOI. The solution predicted a periapsis altitude of 110 km, significantly below the estimated altitude of 150 km. Shortly after the time the spacecraft was expected to emerge from behind Mars, a new orbit, using tracking data to within an hour of MOI, predicted a periapsis altitude of 57 km. At a periapsis of 57 km, the spacecraft would have entered the upper portion of the Mars atmosphere. An analysis by LMA indicated that the spacecraft would have no chance of survival at that low an altitude.

On September 24, 1999, JPL appointed a team to help the MCO Project Team find the reason that the spacecraft signal was lost during the Mars encounter. The team was called the MCO Peer Review Team and was chaired by Dr. Frank Jordan, the former Manager of the Navigation Section and Systems Division (which includes the Navigation Section) at JPL. In addition to Dr. Jordan, the Team included:

- Norm Haynes, former manager of the Systems Division and Mariner 4 Navigation Chief at JPL.
- Dave Smith, former supervisor of the Trajectory Group and former manager of the Mission Design Section at JPL.
- Jim Campbell, former group supervisor in the Tracking Systems and Applications Section and Navigation Chief for the Voyager Project.
- Dave Curkendall, former manager of the Navigation Section.
- Charley Kohlhase, JPL-retired, former mission analysis and Engineering Manager, Voyager, and science and mission design manager, Cassini.
- Fran Sturms, JPL-retired, former supervisor of the Trajectory Group.

- Bill Kirhofer, JPL-retired, former navigation chief for the Ranger, Pioneer, and Galileo missions.
- John McKinney, Telecommunications and Mission Operations Directorate, JPL.
- Glenn Cunningham, JPL-retired, former manager of the Mars Surveyor Operations Project (MSOP).
- Dave Spencer, JPL, mission design and navigation manager, Mars 2001 Project, and former member of the Mars Pathfinder Navigation Team.

The MCO Peer Review Team conducted a fault tree analysis to determine the cause of the MCO loss of signal. Possible causes included:

- Failure of the spacecraft while it was occulted from view.
- Failure of the Deep Space Network (DSN).
- The possibility that the spacecraft had entered the Martian atmosphere and been destroyed.

Analysis of all tracking and telemetry data before entry into Earth occultation indicated that both the spacecraft and DSN were performing as planned. Other missions were being tracked by other DSN stations and reported no problems with the DSN. The final orbit determinations indicated that the spacecraft would pass the planet at a closest approach altitude of just 57 km, far below the target altitude of 226 km and low enough that the spacecraft would enter the planet atmosphere and be destroyed. The time of entry into Earth occultation was 49 seconds early. An analysis of the trajectory indicated that the very early entry into occultation was consistent with a periapsis altitude of about 60 km and independently verified that the spacecraft had entered the Martian atmosphere. The MCO Peer Review Team concluded, therefore, that the inability to reacquire the signal after the expected emergence from occultation was because the spacecraft entered the Martian atmosphere.

Evidence indicates, therefore, that the spacecraft was on a trajectory that was 169 km lower than designed or anticipated. Possible causes were:

- Poor targeting of the TCM-4.
- The spacecraft was deflected off the design trajectory after TCM-4 by, for example, a large gas leak.
- An incorrect estimate of the trajectory before TCM-4.

Analysis of the Doppler and telemetry data before and after TCM-4 indicated that TCM-4 had been conducted as designed and within expected tolerances. No evidence was seen in either the spacecraft telemetry or the tracking data that any kind of unexpected acceleration occurred on the spacecraft between TCM-4 and periapsis. The conclusion of the MCO Peer Review Team was that the cause was a poor estimate of the spacecraft trajectory prior to TCM-4.

The trajectory offset of 169 km was caused by substantial errors one of the following:

- The Mars ephemeris.
- The estimate of the spacecraft trajectory relative to Earth.

The normal expected uncertainties in the Martian ephemeris are about 2 to 3 km. The estimate of the error in the Earth-relative trajectory was approximately 50 to 60 km. To verify the ephemeris, the Mars Global Surveyor ephemeris used at the MGS encounter with Mars in 1997 was substituted for the MCO ephemeris. The result was a 2-km offset, far too small to have caused the large offset in the MCO trajectory. The MCO Peer Review Team concluded, therefore, that failure to acquire a signal from the spacecraft after its expected emergence from occultation was the result of being too deep in the Martian atmosphere, where the spacecraft was destroyed due to aerodynamic forces exerted upon it.

There were several possible causes for the estimate of the spacecraft trajectory relative to Earth to have been in error:

- The radio observables (range and Doppler) used to estimate the trajectory may have been in error.
- The software used to estimate the trajectory from the Doppler and range data may have had a problem.
- There may have been a human input error in the software.
- Some of the supporting parameters used to calculate the spacecraft trajectory — for example, the gravitational constant of the Sun or Mars — may have been incorrect.
- The data filtering method may have been flawed.
- There might have been a spacecraft force modeling error.

The MCO Peer Review Team looked at the radio observables and found that there had been no DSN anomalies reported by MCO or any other of the many missions using the DSN. An alternate, independent version of the software used to estimate the trajectory was used, and resulted in the same estimate of the trajectory. The inputs were scanned by several independent analysts who found no errors. The parameters used in the trajectory-estimation process are contained on a large, carefully controlled lockfile. The lockfile was checked against the one used for MGS on a continuing basis and was found to be consistent. The data filtering techniques were reviewed. The solutions for range only, Doppler only, and range and Doppler were moderately consistent, varying from each other by 50 to 75 km, but were all significantly less than the actual 169-km estimation error. Finally, the MCO Peer Review Team concluded that the modeling of small, non-gravitational forces must be the problem.

The spacecraft is subject to many small, non-gravitational forces, such as the pressure of sunlight, gas leaks from the spacecraft, and Angular Momentum Desaturation (AMD) maneuvers, all of which must be accurately modeled to allow a precision trajectory estimate. The MCO Peer Review Team reviewed the solar pressure model and found no inconsistencies. Further, an error in modeling solar pressure would have resulted in an error perpendicular to the actual error. The MCO Peer Review Team concluded that the solar pressure model was not the problem.

For a leak in the spacecraft to have caused the trajectory offset seen would have required the expenditure of more than a kilogram of gas, which would have been strongly evident in the telemetry of the gas system parameters. No such change was seen.

D 18441

The MCO Peer Review Team reviewed the AMD modeling and found that, early in the cruise phase (approximately April 1999), there was a significant offset between the expected values of velocity increment along the line of sight calculated from the AMD files and the value obtained from the Doppler data. The MCO Peer Review Team learned that the velocity increment to the spacecraft from about two AMDs per day was more than a cm/second per day and would result in a trajectory error of several hundreds kilometers unless accurately modeled.

On September 27, 1999, the MCO Peer Review Team requested a teleconference with the LMA Attitude Control Team and the Navigation Team to discuss possible AMD modeling problems. The teleconference took place on September 28, 1999, and the MCO Peer Review Team asked LMA for help in diagnosing the problem. The next day, September 29, 1999, LMA called to report that they had delivered the impulse applied to the spacecraft in pounds-force seconds, instead of Newton-seconds, as specified in the interface agreements. The MCO Peer Review Team asked the Navigation Team to run a new trajectory estimate from before TCM-4, keeping everything the same except the impulse delivered to the spacecraft, which was multiplied by 4.45, the conversion factor between pounds-force and Newtons. The new estimate yielded a trajectory 160 km closer to Mars. If TCM-4 had been based on this trajectory, the periapsis altitude would have been 188 km rather than 57 km and the spacecraft would not have entered the Martian atmosphere.

Figure 1 is a fault-tree analysis of the MCO Peer Review Team's investigation.

Appendix 2. Navigation and Small Forces Modeling

A. Navigation Process

The navigation process consists of estimating the trajectory of the spacecraft and, if necessary, performing Trajectory Correction Maneuvers (TCMs) required for mission success. The major effort for this process involves the estimation of the trajectory. The computation of the corrective maneuvers is generally a less demanding task. The trajectory estimation process is commonly referred to as "orbit determination."

B. Forces Affecting the Motion

The spacecraft trajectory is dependent upon all of the forces that act upon the spacecraft. The gravitational attraction from the Sun and all of the planets are obvious factors. Another factor is solar pressure from solar photons interacting with the spacecraft surface.

A spacecraft can even exert forces upon its own motion by the release of gases that either leak out or that result from the intentional firing of thrusters to alter the trajectory. In short, any material ejected from a spacecraft will alter the spacecraft's motion path by producing a force in the opposite direction of the ejected mass.

A spacecraft has the periodic need to alter its attitude orientation, either to perform TCMs in a particular direction, point downlink telemetry antennas, or collect scientific data. This may be achieved by the firing of attitude-control thrusters or by the use of reaction wheels (also called momentum wheels) to make turns about the spacecraft's three (typically) main axes of rotation. Reaction wheels are like spinning gyroscopes whose speeds can be altered to cause the spacecraft to turn about one or more of its axes. From time to time, reaction wheel speeds may need to be brought back within certain operating bounds, and attitude-control thrusters are fired to allow "desaturation" of the wheels. In the parlance of the MCO, these attitude-control thruster firings are referred to as Angular Momentum Desaturation (AMD) events.

The attitude-control thrusters may be "coupled" or "uncoupled." In a coupled design, thrusters are mounted on the spacecraft in such a manner that the firing of a particular thruster to rotate the spacecraft is always offset by another thruster pointed in the opposite direction, thereby resulting in no net thruster gas ejected in any one particular direction. In an uncoupled design, however, the thrusters are pointed in the same direction and impart a net force on the spacecraft that will both alter its trajectory and achieve the desired attitude change. Any small forces that act on the spacecraft in this manner will alter its path of motion by changing its velocity; deep-space navigators often refer to these tiny velocity increments as "delta-Vs."

Orbit Determination Measurements

Estimating the orbit of a spacecraft is done by collecting many measurements of the spacecraft's position and motion in order to "solve for" the most likely path that it is on now, and then predicting where it will be in the future. We can track the spacecraft radio signal using large radio telescopes on Earth, or we can use suitable spacecraft cameras to transmit photographs of the destination target against a background of known stars. However, because our measurements are never absolutely precise, we can never know the spacecraft orbit with absolute precision. On the other hand, we have the capability to collect thousands of measurements over many days, process them in a sophisticated Orbit Determination Program (ODP), and develop a pretty good estimate for the spacecraft orbit.

The primary measurement types we get between the tracking antennas on Earth and the spacecraft are 1) the range and 2) the rate at which the range is changing, known as the Doppler. The range can be measured to an accuracy of a few meters; the Doppler can be measured to better than 0.1 mm/s. Both of these measurements are, however, only available along the line-of-sight (LOS) to the spacecraft. We can only "pick up" what is happening at right angles to the LOS by allowing enough time to pass for the spacecraft to move some distance along its curving orbit. If one tries to "re-determine" the spacecraft's orbit after the execution of a TCM, it may not be possible to do so quickly without taking advantage of what was known before the trajectory correction was made. Therefore, the degree of conviction that is placed upon that prior knowledge must be done with care.

Orbit Determination Process

As you might expect, the ODP must know a great deal, not only about all possible forces that may be acting upon the spacecraft, but also about every factor that could potentially alter the value of one of the important measurements. For example, the ODP must know within a few centimeters the exact locations of the large tracking stations on Earth. The ODP must also account for many other considerations, such as the medium through which the radio signal must travel between Earth and the spacecraft, the velocity of light, the spacecraft communications equipment, etc. JPL's high-precision ODP is able to estimate several hundred different quantities that can influence either the motion of the trajectory or the value and calibration accuracy of the key measurements that are being collected and processed.

Internally, the ODP makes small adjustments to many of these "model parameters" in order that the errors seen between the actual measurements and the calculated measurements (i.e., the measurement values that would have resulted if the spacecraft were following a certain orbit) are as small as possible. However, each orbit solution depends upon how much data and what types are used, what parameters are estimated, what the original accuracy ("a priori" knowledge) of each of these parameters is believed to be, and other controls available to the navigation experts. In navigation parlance, much of this process is referred to as "covariance analysis."

When the navigators are trying out different combinations of solutions, the ODP is adjusting its estimates for the various solve-for parameters within the bounds established by the a priori uncertainties originally assigned to each of these parameters. If, however, the true error in one or more of these parameters is much larger than assumed by the choice of a priori error bounds, the ODP will adjust some other parameters erroneously in its mathematical quest to achieve the lowest possible residuals. In fact, it may adjust its own best estimate of the spacecraft position and velocity (known as spacecraft "state") in seeking to reduce the residuals. If this is further compounded by using the new solve-for knowledge to forecast where the spacecraft will be in the future, a growing error can result. Therefore, our initial knowledge of many key parameters must be carefully and accurately chosen.

The Navigation Art

Most of the time, the navigation process goes smoothly and each deep-space mission is guided with high precision to its destination in a manner that supports the scientific objectives of the mission. Occasionally, however, orbit solutions do not "converge" to a single place, but, instead, scatter about and indicate that something is wrong. Clearly, the problem lies somewhere in the complex model of all the forces and all the measurements, and must be investigated for navigation success or to alert project management that greater tolerance must be allowed for the effect of possibly larger-than-expected navigation errors. The investigation process requires that both the art and science of the navigation trade be applied with great skill.

In the case of MCO, the model used to represent the small delta-Vs resulting from reaction wheel desaturation activities (occurring some dozen times each week) was not accurate. The errors were sufficient to result in a trajectory displacement of approximately 169 km in where the spacecraft would pass Mars at its closest approach. Unfortunately, the consequences of this total error were in the direction closer to Mars, causing the spacecraft to take too deep a cut through the Martian atmosphere and, thereby, fail to survive. The unrecognized error had been building up for months along the spacecraft y-axis direction, perpendicular to the tracking LOS and the spacecraft's plane of motion, and was, therefore, not easy to see in the data. Had the model of these small forces been correct to begin with, or had the navigation process detected and corrected these model errors in flight, or had larger navigation uncertainties been used in targeting the final aiming point at Mars, the MCO failure would not have occurred.

C. Modeling of Small Forces from Spacecraft Thrusters

In order to describe the process of modeling — i.e., estimating the magnitudes and directions of the small forces produced by the spacecraft thrusters — the Mars Global Surveyor (MGS) will be used as an example. Since the MCO's approach was based on the MGS approach, this example will also describe the MCO's heritage base. As was pointed out above, reaction wheels accumulate a net spin rate, due primarily during cruise to the action of unbalanced solar pressure torques on the spacecraft. As a result, the reaction wheels have to be desaturated periodically. For the MGS, the unbalanced thrusting that occurs at these times is estimated by ground software using certain spacecraft telemetry data. Thruster on-time data, spacecraft attitude and spin rate, plus propellant tank pressure are processed using equations provided by the thruster manufacturer. The output consists of the net impulse (Newton-seconds) delivered by each thruster, plus the time and spacecraft orientation at which the impulses were imparted to the spacecraft. Navigation software then combines the individual thruster impulses into a single vector impulse and divides by the spacecraft mass to obtain the delta-V. The delta-Vs from all previously performed desaturations are applied during the OD process to both estimate future desaturation delta-Vs and estimate the current and future spacecraft state.

On the MGS, the accuracy of the desaturation delta-Vs calculated directly from the spacecraft data was 80 to 90%, which is considered good. Desaturation delta-Vs based on pre-launch estimates were predicted to average ~9 mm/second per week during cruise. This consisted of ~1 mm/second from yaw unloads performed every 5 days and ~14 mm/second from spin unloads performed every 14 days. Yaw unload delta-Vs were small because the spacecraft nominally spun around the Earth line (0.01 rpm), allowing the unloading of one component of momentum to be timed such that it was accomplished using coupled thrusters. The remaining component of momentum was unloaded during the spin unloads using uncoupled thrusters. It should be noted that while this spin unload delta-V was large, over time it occurred at random spin angles, so errors in its estimated size tended to cancel out.

D. Comparison of MGS with the Mars Observer

The Mars Observer (MO) desaturation delta-Vs tended to be smaller (~1 mm/second for all events) and less frequent (~ once per week). Although the MO and the MGS used the same reaction wheels and both spun around the Earth line during cruise, the MO could use coupled thrusters for all desaturation events. In addition, the MO had a more favorable solar array configuration. The MO array was configured symmetrically during cruise, which resulted in low solar pressure torque applied to the spacecraft. The MGS array consisted of two wings that were canted to each other during cruise. Although they were nominally canted symmetrically about the Earth line, they were not symmetric about the Sun line. Thus, the evolution from the MO to the MGS spacecraft design resulted in the potential for more difficulties in accurately modeling spacecraft small forces. Nevertheless, since the MGS desaturation delta-Vs were reported with good accuracy, and since the delta-V errors for the uncoupled thrusting tended to self cancel, MGS's less desirable features had little impact on OD accuracy.

E. Evolution to MCO Small Forces Modeling

After the MGS, the MCO spacecraft was developed along with the MPL spacecraft by LMA, which was concurrently developing the Stardust spacecraft. While the MCO spacecraft and mission were basically similar to the MGS, the MPL and Stardust were very different. This led to more emphasis and a new approach to small forces estimation for the latter two spacecraft.

Stardust put more emphasis on small forces modeling during development due to their unique sensitivity. Because of their thruster orientations and seven-year mission, the cumulative small forces delta-V was estimated at 20m/second. The propellant margin would not have been sufficient to compensate for the delta-V effect after it was observed. Therefore, the delta-V had to be predicted accurately and compensated for early in flight. Consequently, during mission operations, in-flight calibrations were performed and changes were made to improve the modeling accuracy beyond the initially observed accuracy of 90%.

Neither the MPL nor Stardust uses reaction wheels, and both have very tight targeting requirements. Furthermore, neither employ coupled thrusters. Delta-Vs are imparted to the spacecraft during normal attitude control "deadbanding." These delta-Vs are smaller and more frequent than those associated with reaction wheel desaturations. In order to handle the small forces information from the spacecraft more efficiently, the Stardust project negotiated a new interface for this data. This consists of the spacecraft flight software processing thruster firing data onboard and providing small forces delta-V (m/second) directly via a "small forces packet" in downlink telemetry. Navigation software uses the delta-V data obtained from this packet directly.

LMA baselined the new approach for the MCO and MPL spacecraft because it was efficient to use common software modules on all three spacecraft. The Mars '98 project agreed to this approach, in particular because it made sense for the MPL. Later, when the MCO Navigation Team Chief transitioned onto the MCO from the MGS, the new approach was questioned for the MCO. It was argued that, since the MCO was similar to MGS and the small forces software worked fine on the MGS, there was no compelling reason to change. Furthermore, new software on the navigation side of the interface would be required for handling the new data content for the MCO. Instead, the project decided to develop an MGS-like interface for MCO (see the Background section of Appendix 3).

F. Comparison of MCO with MGS

There were two significant differences between the MCO and the MGS with regard to small forces:

- The MCO had a single solar-array (S/A) wing, offset a large distance from the spacecraft center of mass
- The MCO did not spin during cruise.

The large S/A wing offset resulted in a large solar pressure torque in the cruise configuration. This torque significantly increased the amount of thrusting required to desaturate the two reaction wheels that had to oppose it. One of these two wheels was desaturated by partially coupled thrusters. The other wheel was desaturated by uncoupled thrusters having most of their resultant delta-V acting in the spacecraft –Y direction. (Desaturation of either wheel resulted in a small component of thrust in the spacecraft +Z direction.)

Unlike the MGS, the MCO did not spin during cruise. However, the baseline at the spacecraft Critical Design Review was to rotate the spacecraft 180 degree about the Sun line up to twice a day, thereby canceling much of the reaction-wheel momentum buildup and reducing thruster-use frequency. The purpose of this was to conserve propellant. This strategy would have practically eliminated the uncoupled desaturation thrusting that resulted in the –Y direction referred to above. The 180 degree rotation strategy was later abandoned, however, because it resulted in rotating the high-gain antenna off the Earth line. Off pointing of the antenna complicated sequencing with respect to downlink pass scheduling, and the propellant savings was deemed not significant.

During inner cruise (i.e., prior to April 1999), the spacecraft –Y direction was toward the trajectory north pole, while during outer cruise, the –Y direction was pointed toward the trajectory south pole. The spacecraft went through inferior conjunction about April 8, 1999, and the relative positions of the Earth and Sun reversed. The effect of this reversal was a partial cancellation of the –Y delta-V over time. To calculate or even approximate the amount of this cancellation is difficult. The solar pressure torque about spacecraft +X has to be integrated over the mission, and the torque varies with time as a function of solar range, Sun-probe-Earth angle, solar array offpoint angle, and spacecraft antenna offpoint angle. The fact that the spacecraft at Mars Orbit Insertion (MOI) was closer to the planet than estimated indicates that –Y delta-V during inner cruise was not dominant over that during outer cruise. This does not seem unreasonable based on an examination of the relative duration of the two cruise phases and the smaller offpoint angle of the array during outer cruise.

Appendix 3. Small Forces Modeling and Format Errors

A. Background

The LMA Spacecraft Development Team had wanted to use the same small forces interface for the MCO that was used for the Stardust and Mars Polar Lander programs. That interface used the delta-V calculation performed by flight software and placed in the telemetry data stream.

The Spacecraft Development Team decided to essentially eliminate the AMD file from the Spacecraft Performance and Analysis Software (SPAS) software development for the MCO, the MPL, and Stardust. The approach was to have the delta-V calculations done on board the spacecraft and embedded in the engineering telemetry stream. When the Mars Surveyor Operations Project (MSOP) Navigation Team began to be involved with the MCO (approximately a year before launch), they objected to this approach because it would mean:

- Changing operational software at JPL that was inherited from and had worked well for the MGS.
- The MCO Navigation Team would have to access the spacecraft telemetry data and find the right parameters.

The Navigation Team preferred using the same AMD file format that was used successfully on MGS and requested that the AMD file have exactly the same format as defined in the MGS Spacecraft Interface Specification (SIS). To do this, LMA would have to develop a program that would take the new small forces packet, discard the delta-V data, and convert the remainder of the packet's data, which included spacecraft-thruster on time, into the same data as is provided by the MGS ground software program (AMDGEN). The new program was projected to have high inheritance from AMDGEN, requiring approximately 10% new code. New equations would be required for converting thruster on time into impulse, because the MCO and MGS thrusters are different.

LMA preferred the onboard approach because it was already in place, because it maintained commonality with MPL and Stardust, and because the alternative would require having to do some curious things, like filling in all 12 thruster packets (even though MCO only had 8 thrusters). In the end, the change requested by the Navigation Team was made without assessing the relative effort of modifying the MGS SPAS software at LMA versus modifying the MGO navigation software interface at JPL.

B. Small Forces Modeling Error

The MGS Software Interface Specification (Angular Momentum Desaturation File Software Interface Specification (SIS) EAE003) was cited for the interface requirements. That interface specification required the impulse bit (IBIT) units to be in Newton-seconds.

A Software Requirements and Design Specification (SRDS) was written in preparation to developing the Sm_forces software program for generating the AMD files requested by the Navigation Team. The Sm_forces SRDS did not directly specify the units for any of the output parameters. However, the file format section of this document had a sentence at the end of the section that said, "For further details on the Output file format, refer to SIS EAE-003."

SIS EAE003 required the IBIT to be in Newton-seconds. The equation for the calculation of the IBIT was not in the SRDS or in any of the other software design documents. The equation was provided on a page extracted from another document (provided by Primex, the thruster contractor) and placed in the software development folder. The description of units on that page shows that the units of the IBIT calculation are pounds-force-seconds. The walkthroughs done on design, code, and test did not catch the error in the units. In addition, not all of the people required to attend the walkthrough were there. Finally, walkthrough training is no longer provided for the software people.

The test case for the unit test did not detect the error. Most likely, the predicted output for the test was determined from the equation provided by Primex. This equation produces the answer in pound-seconds. A formal interface test between the small forces software and the navigation software was not performed. Had even a simple interface test been performed, the format problems with the interface would have been discovered before launch. Depending on the extent of the interface test, the units problem still may not have been discovered. However, interface testing was not performed, which violated the Flight Systems Software Management and Development Plan.

An Operational Readiness Test (ORT) for TCM-1 was performed prior to the launch of the Mars Climate Orbiter. Also, the Spacecraft Team ran the Sm_forces program to produce an AMD file for the Navigation Team to use during the ORT. The ORT was declared successful without having confirmed that the interface file was read successfully. Non-conformance documentation was not generated during or after the ORT. The Review Board was told that the Navigation Team was not aware that the small forces output file was on the server at the time of the ORT and, as a result, the Navigation Team did not attempt to read the file.

The spacecraft launched with the error in units still in the system.

Figure 12 is a flow diagram of the events described in this section.

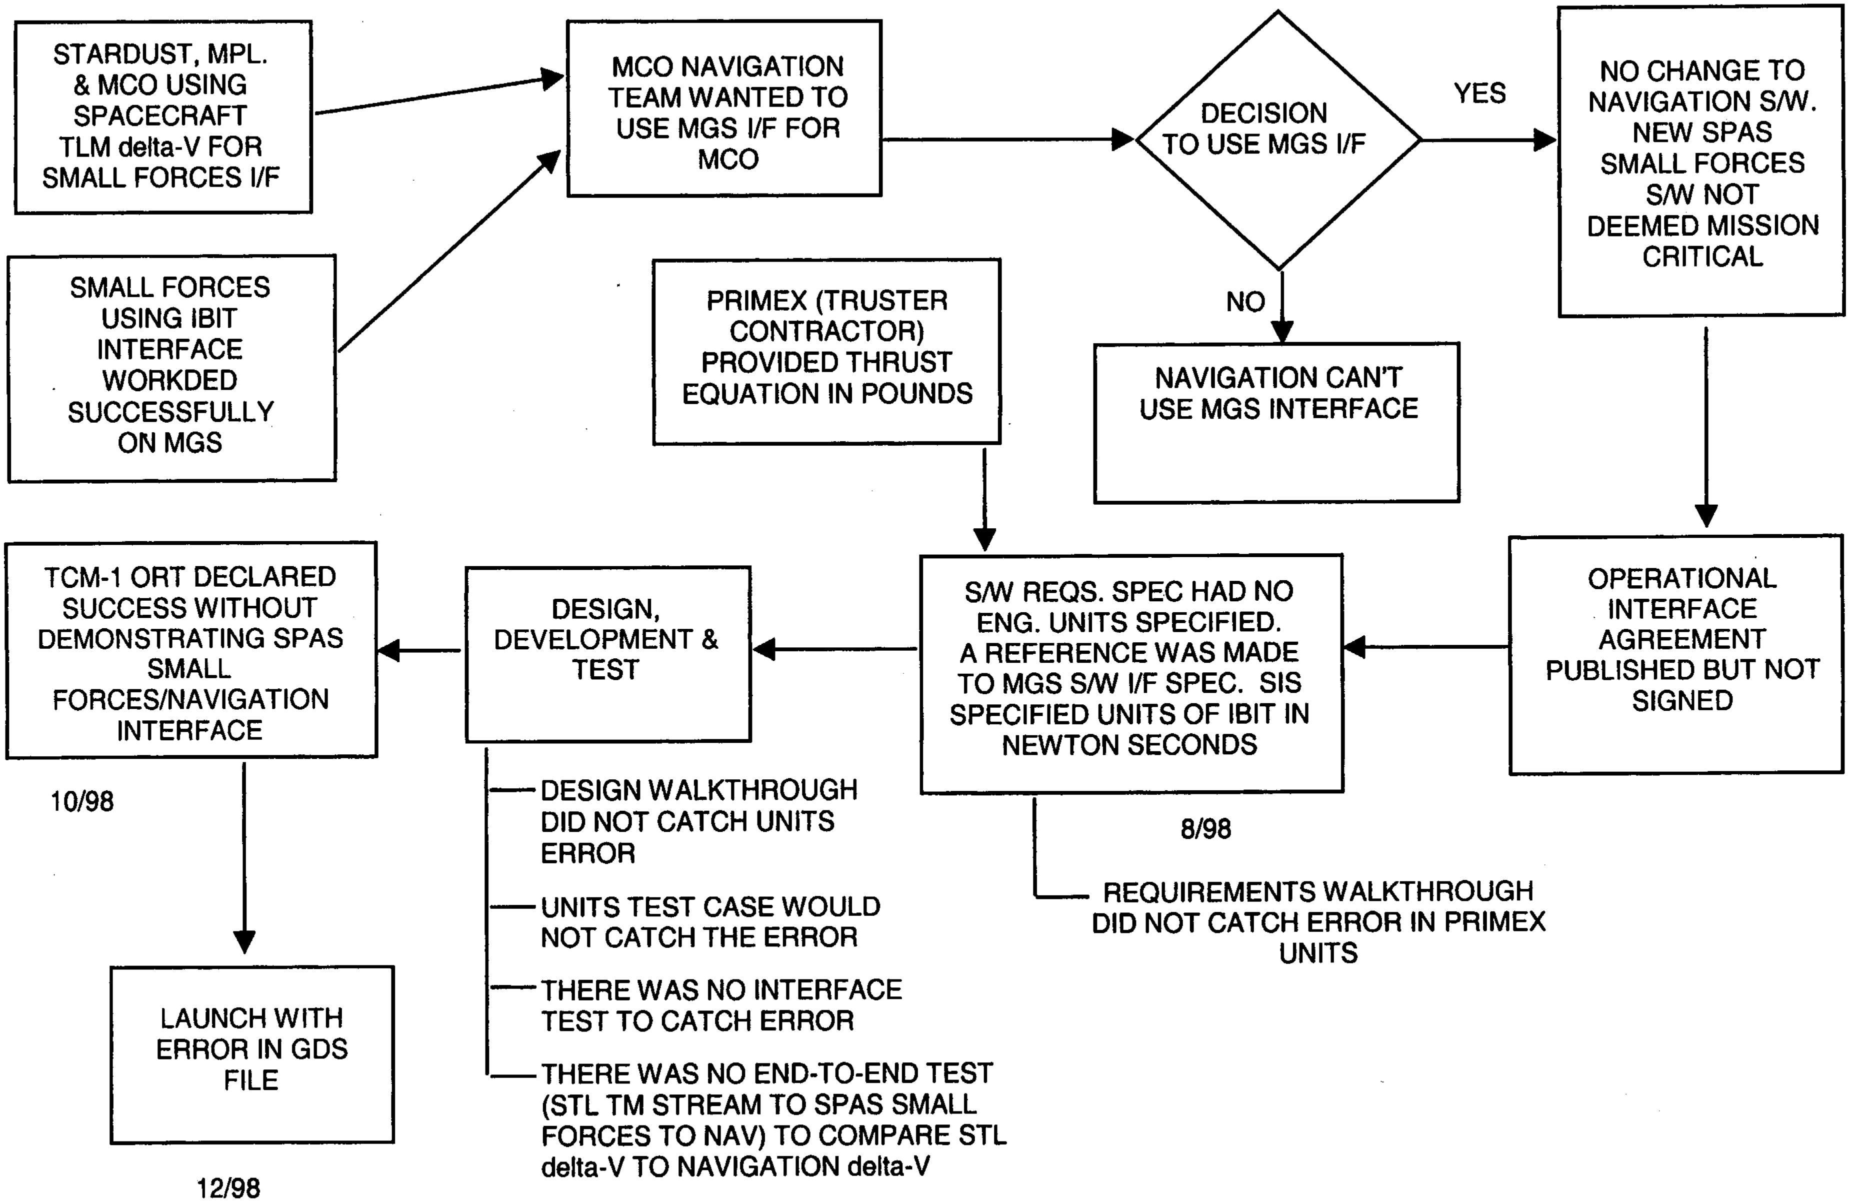

Figure 12: Units Error Events

C. Format Error

Within two days of launch there were two major discoveries regarding the AMD files:

- The files transmitted from LMA to JPL were in a format inconsistent with the navigation software.
- The coordinate system for the quaternions was MME instead of EME.

At that time, it was also noted that there were 2 desaturations per day. Although the number of desaturations was in accordance with the predicted spacecraft performance, the number of desaturations was a complete surprise to the Navigation Team. In addition, there were approximately 60 AMD files per desaturation maneuver. It is clear there was no valid interface testing for the AMD files during the launch readiness test. The 60 or so files resulted from the desaturation maneuver strategy. Rather than one long firing, a short burn was implemented during each computer cycle. Each short firing was encapsulated into a separate file.

The file formats were not compliant with the requirements in the SIS. In addition, the SIS had not been kept up to date and was inconsistent with the navigation software interface that was currently in use on MGS. These two issues were described in an e-mail from JPL to LMA on December 20, 1998. The Navigation Team requested that LMA fix the format to make it consistent with the navigation software interface requirements. Navigation also requested that the files be consolidated into one AMD file per desaturation maneuver. At LMA, an analyst was given the task to consolidate the individual pulse files into one file per desaturation. In the January 1999 to February 1999 time frame, the analyst was also working Stardust attitude issues that were perceived to have higher priority.

Knowing the Stardust attitude issues, JPL requested that, at least, the times of the desaturations events be provided. Knowing the times of the desaturation events, even if the magnitudes and directions were not known, would be an aid to the Navigation Team's ability to estimate the small forces. The desaturation times were provided by LMA on February 20, 1999. Also during the January 1999 to February 1999 time frame, the Navigation Team apprised the project staff that the small forces problem could affect OD accuracy. By February 11, 1999, a navigation analyst at JPL had manually manipulated the AMD format to produce a readable file. Results were faxed to LMA on March 12, 1999, showing the effect of a corrected AMD file on the two-way Doppler data. On April 6, 1999, a Failure Report (FR) was written at LMA describing the file issues. On April 19, 1999, the file was verified as corrected at LMA. An e-mail from Navigation on April 26, 1999, verified that the format was readable. In fact only the format issue had been resolved. The units error was still present and unrecognized.

Figure 13 outlines the format error.

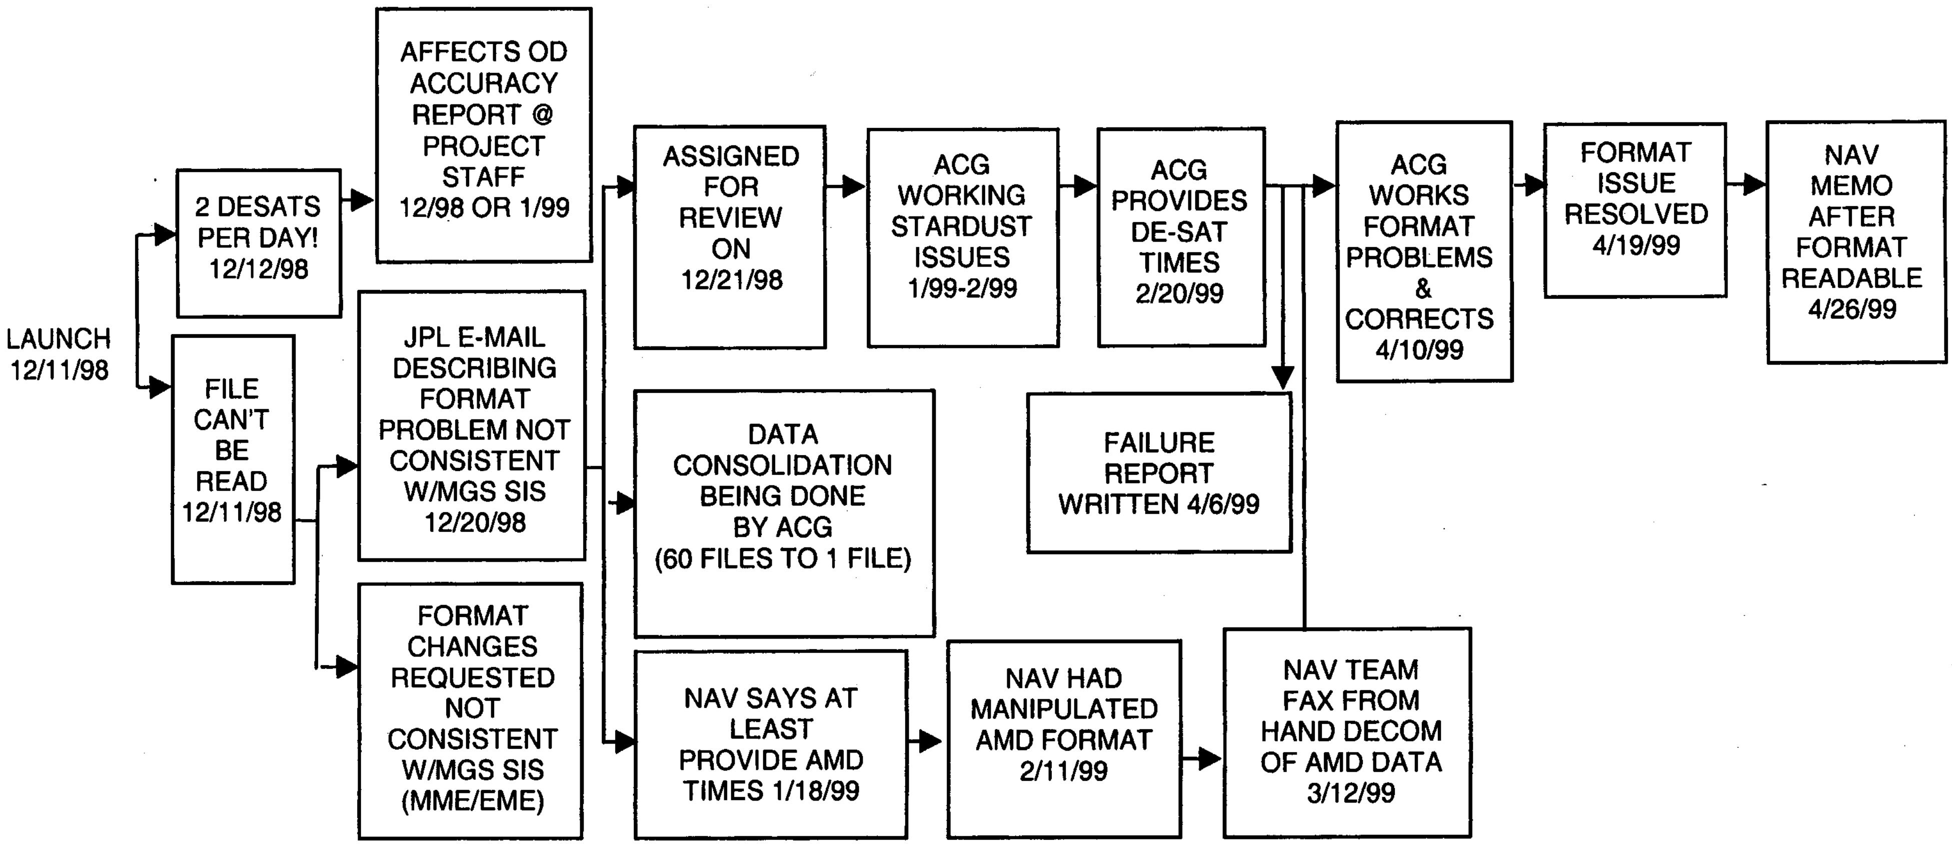

Figure 13: AMD Format Problem

Appendix 4. In-Flight Handling of Small Forces Error

A. Background

The MCO and the Mars Polar Lander (MPL) were a part of the Mars '98 Development Project. The Mars '98 Development Project was responsible for the development of the MCO and the MPL, turning responsibility for the operations over to the Mars Surveyor Operations Project (MSOP) at launch. The MSOP is responsible for the operation of all Mars Surveyor missions, which, at the time of launch, included the Mars Global Surveyor (MGS), the MCO, and the MPL.

During development, navigation studies were undertaken to determine with what accuracy the MCO and MPL could be navigated to their target orbits or landing sites. These studies were done by the Development Project with no participation by the MSOP Navigation Team, which would be navigating the missions after launch.

The preliminary studies for the MCO were based on a spacecraft design that produced one AMD maneuver every 5 to 7 days. Later, the Project considered an attitude implementation called the "barbecue" mode that would have reduced the number of AMD maneuvers significantly in order to save propellant mass. This strategy consisted of rotating the spacecraft 180° about the Sun line twice per day.

Eventually, the project decided against the barbecue mode option on the basis of required operational complexities. Navigational inputs were not considered in coming to this decision. This strategy, which would have been beneficial to navigation, was abandoned because of communication-link sequencing complications and because the propellant savings was not significant. The negative impact on navigation was not considered, nor was the Navigation Team informed of the change.

Prior to Preliminary Design Review, the project chose an attitude-control implementation mode that would result in about 2 AMD maneuvers per day, an increase of more than a factor of 10 in AMD maneuver frequency. The navigation accuracy analysis should have been done to determine the impact of this increase. If a realistic navigation study had been done by the MSOP Navigation Team using 2 AMD maneuvers per day, and if the findings of the study had been presented at the decision meeting, a different decision might have been made regarding the cruise attitude control scheme. Barring this, the Navigation Team would have at least been aware of the criticality of the AMD maneuvers.

The desaturation delta-V was large due to large solar pressure disturbance torque caused by the asymmetrical location of the solar array, and because of the use of unbalanced thrusters to unload the reaction wheels. However, the cumulative desaturation delta-V would have been significantly reduced if the "barbecue" strategy baselined at the spacecraft Attitude Control System Critical Design Review had been used.

In addition, no estimates of cumulative desaturation delta-V magnitude and direction was provided, even though for cruise they were largely deterministic. Because the magnitude was large, the deterministic nature of the estimates should have been considered in the navigation plan and provided to the flight Navigation Team. A better understanding of the spacecraft would have led to an understanding that the desaturation delta-Vs should not be approached as statistical. The real effects were masked or minimized by considering them as statistical in nature.

Shortly after launch, in the December to January timeframe, the Navigation Team was surprised at the frequency of the AMD maneuvers. In flight, the maneuvers were occurring about twice per day, instead of once every 5 to 7 days, as the pre-launch studies of navigation accuracy had assumed. The Navigation Team raised the issue at an MSOP staff meeting and questioned whether they could achieve the preflight navigation accuracy that had been predicted with the great increase in AMD maneuvers. Navigation also raised the issue of the large increase in AMD maneuvers with the Spacecraft Team and was told that the Spacecraft Team expected a frequency of about 2 a day with the attitude control design that had been agreed on well before launch.

Apparently, there had been no information on the increased frequency of the AMD maneuvers transferred from the Development Project, which conducted the pre-launch navigation-accuracy analysis, to the Navigation Team that would conduct the actual in-flight navigation analysis. In addition, the Development Project had not redone the navigation accuracy analysis with the 2 per day frequency rather than the once per 5 to 7 days they had assumed in the original study.

B. Handling of the Delta-V Discrepancies

On April 14, the Navigation Team Leader visited LMA to provide a tutorial on non-gravitational forces and emphasized the importance of non-gravitational force modeling on trajectory reconstruction. He demonstrated the effects of the AMD maneuvers on the MCO Doppler residuals.

Two days after the tutorial, the Spacecraft Team sent an e-mail to the Navigation Team thanking them for the tutorial and re-iterating that the SIS was the controlling document for the AMD file interface.

On April 26, the Navigation Team sent an e-mail, with copies to the Spacecraft Team thanking them for fixing the file format problem and indicating that the Navigation Team could now read the AMD data that was being sent to them. The e-mail also indicated that in several AMDs that had occurred while the spacecraft was being tracked, the Doppler measured value of the AMD component along the Earth direction was several times larger than the value from the AMD file. The Navigation Team asked the Spacecraft Team to look into the problem. No ISA was written at JPL or at LMA.

After receiving the Navigation Team memo, the Spacecraft Team initiated the investigation into the discrepancy in AMD delta-V magnitude. The Spacecraft Team, on May 6, 1999, sent an e-mail to the Navigation Team indicating that the delta-V discrepancies would be investigated and requested more data. The Spacecraft Team indicated that it would compare the value of the Doppler-derived delta-V with the value calculated in the flight software and sent to the ground in the telemetry stream. On May 10, 1999, the Navigation Team delivered plots of the data to the Spacecraft Team. On May 12, 1999, the Navigation Team Leader visited LMA to conduct aerobraking discussions regarding the MCO. While the Navigation Team Leader was at LMA, he had a side meeting with some of the Spacecraft Team members to discuss the velocity differences expressed in the Navigation Team memo. On May 14, the Spacecraft Team sent an e-mail to the Navigation Team acknowledging the receipt of the plots and asking for the data in tabular form and in spacecraft event time. On May 21, the Navigation Team sent the requested data in tabular form. This was the last written communication between JPL and LMA on this subject until after the mission failure. Subsequently, the Navigation Team made a couple of telephone calls to the Spacecraft Team regarding the status of the investigation and was told they were still working on it. In this time frame, the Navigation Team Leader asked project management to escalate the priority of the AMD delta-V magnitude investigation at LMA.

Project management declined to escalate the priority. Instead, project management told the Navigation Team:

- That the Navigation Team needed to do the best they could with the data they had.
- That what the project really needed was the Navigation Team's best estimate of the uncertainty in the trajectory.

On June 9, 1999, the Attitude Control Group (ACG) at LMA was reorganized. The ACG previously had a single point of contact for the group with all projects. In the reorganization, separate persons were established as the points of contact for each project. A person not previously working on MCO was assigned to be the MCO point of contact. The analyst working on the delta-V problem was assigned as point of contact to Stardust, but also retained the MCO AMD task. The new MCO point of contact was apparently not told of the MCO AMD task.

In June, the estimates of the trajectory following TCM-2 had begun to change significantly. Believing that the velocity discrepancy investigation would yield no new information because they had heard nothing from LMA, the Navigation Team began to solve for the components of the AMD maneuvers. The Team used the AMD file information as the input and 1 mm/second as the a priori uncertainty in each component.

By early August, after TCM-3, the orbit solutions, projected to periapsis altitude, dropped lower. In addition, the solutions using range-only data, Doppler-only data, and range-plus-Doppler data exhibited a spread far greater than the uncertainty in any of the individual solutions. The Navigation Team informed project management that the periapsis target altitude was too low given the problems they were having with fitting the AMD data and with the scatter in the solutions. By late August, numerous meetings relative to the scatter in the solutions were held between project management, the Navigation Team, and a consultant. However, there was apparently no linking of the scatter in the solutions with the AMD velocity discrepancy and, in fact, the AMD problem may not have even surfaced at these meetings.

On September 9, 1999, at the MOI and Aerobraking Review, the Navigation Team reported no concerns relative to MOI, with only a mention of the AMD velocity problem, and displayed no concerns regarding the orbit-solution scatter.

TCM-4 was conducted on September 15, 1999. On September 16, during the meeting with the project to go over the quick results of the maneuver, the first indication was that this maneuver had also come in low (approximately 60 km lower than expected).

After the meeting, an inquiry was made regarding progress on the AMD velocity problem. It was determined that there was none. There was still no connection made between the results of the maneuver, the scatter in the orbit solutions, and the AMD problem.

Figure 14 illustrated the chronology of the desaturation delta-V anomaly.

April 14
LMA
The Navigation Team provides a tutorial on navigation. Emphasizes importance of impact on non-gravitational forces on trajectory reconstruction. Discussed AMD effects in MCO Doppler data and residuals

April 16
LMA
Spacecraft Team e-mail to Navigation Team confirms that the SIS is the controlling document

April 26
JPL
Memo from Navigation Team shows Doppler line-of-sight different from AMD file (April 16 to April 21). No ISA written.

Apr 26 – May 6
ACG receives AMD file investigation assignment.

May 6
LMA
ACG e-mail – says delta-V differences being investigated; requests more data. Indicates investigation still underway. Will compare to LMA delta-Vs.

Week of May 10
JPL
JPL delivers plots in EME2000 of delta-Vs observed.

The Navigation Team at LMA for aerobraking discussion. Also discusses velocity differences from Navigation Team memo – investigation still in process.

May 12 – LMA
May 14 – LMA
Spacecraft Team e-mail acknowledges plots and indicates investigation is still underway. Requests data in tabular form.

May 21
JPL
Navigation Team delivers tabular forms delta-Vs – 2 tables - 3D.

Late May – Early June
LMA
The Stardust delta-V problem is assigned to the ACG, but they never got back to the MCO delta-V problem.

Late May – Early June
Telephone calls from Navigation Team to Spacecraft Team requesting an investigation report.

Early June
JPL
Trajectory movement post TCM-2 noticed over a 2-month period. Decision made to solve for delta-V. Solutions stabilized.

June
JPL
Navigation Team assumes investigation will yield no new information

June 9
LMA
Reorganization of the ACG. New MCO point of contact assigned.

Date unknown
JPL
Navigation Team begins solution techniques to solve for delta-V using AMD files as input data and 1 mm/set as a priority uncertainty.

June – July
JPL
Navigation Team asks project management to intercede with LMA on delta-V investigation

June – July
JPL
Project management tells Navigation "No compelling reason to target low", "not critical to achieve given level of accuracy". Directs Navigation to "figure out how to do it with the best you have".

Early August
JPL
Post TCM-3 solution goes downward. No recognition that TCM-3 was absorbing large delta-V effect

August - JPL
Navigation Team tells Spacecraft Team that the periapsis alt was too low given the AMD accuracy

Mid August – Early Sept
Numerous meetings held between project management, Navigation, and consultants. No linking of solution spread to delta-V problem. Delta-V solution evidently not shared and discussed.

September 9
JPL
No delta-V concerns shown in MOI readiness review but did not show delta-V plots.

September
JPL
ACG calls Navigation Team. Doesn't pursue investigation

September 16
JPL
Navigation Team questions MCO point of contact regarding delta-V investigation. MCO point of contact's first knowledge.

Post
LMA
TCM-4 quick-look Drop – no connection to delta-V plots.

Figure 14: Chronology of the Desaturation Delta-V Anomaly

Appendix 5. Risk Strategy

A. Planned Orbit Initial Design (August 1997 Mission Plan)

The periapsis altitude on orbit 2 (P2) altitude was lowered from 250 km (used for the MGS) to 180 km for the MCO. The basis for this selection was to reduce MOI delta-V and achieve earlier aerobraking and readiness to support the MPL. The project risk strategy consisted of assuring a 99% probability that P2 would not be lower than 150 km, an altitude for which some aerodynamic heating concerns existed. With an estimated 3-sigma navigation delivery error of roughly 30 km, that led to an initial choice of 180 km for P2. This value could be adjusted in flight if larger navigation errors were forecast. The Navigation Team assumed that the altitude at P1 would be safe as long as the P2 criteria were satisfied.

The August 1997 version of the Mission Plan indicated that TCM-5 could occur as late as MOI — 2 days as a contingency maneuver, if needed. There was no consideration given to whether or not TCM-5's placement that late in the timeline would be of any risk concern to an MOI sequence that would be aboard and "clocking out." No risk analysis was done to allow for any downside threat for larger-than-believed navigation errors. The principal rationale was to ensure that P2 remained above 150 km, using known accuracy estimates. In any event, there was further safety in the knowledge that serious atmospheric disruption to the spacecraft was not likely a problem at periapsis altitudes greater than 120 km.

B. Navigation Targeting Meeting for TCM-4 (9/2/99)

At a prior meeting on August 30, 1999, as a result of concerns regarding the orbit solution spread, the Navigation Team requested a TCM-4 aiming point corresponding to a P2 altitude of 250 km. Project management asked the Navigation Team to run more cases and report the results at the TCM-4 target-selection meeting scheduled for September 2, 1999. At the September 2 meeting, as a result of a Navigation Team's 3-sigma error quote of 60 km, a final P2 altitude of 210 km (P1 of 226 km) was chosen by project management. To allay concerns, the Navigation Team was told that TCM-5 could always be used to raise the orbit if such an action appeared warranted. This was in accordance with the project risk strategy.

Furthermore, there had been analysis and discussion of the possibility of designing an inefficient MOI maneuver for the final MOI sequence that could raise P2 altitude by an additional 20 km. This possibility had been assessed between September 7 and September 17, 1999, but LMA personnel were never requested to build or test such a sequence.

The Navigation Team again requested that project management further raise P2 to 250 km on September 13, because of the volatility of OD solutions. This time Navigation acquiesced to not changing the targeting on the condition that there would be a TCM-5 and that the MOI maneuver would be modified as mentioned above. There was a misunderstanding between Navigation and project management regarding whether this condition was actually agreed to.

C. Post-TCM-4 Activities Relating to TCM-5 (9/15-9/20/99)

Based upon 2 hours of post-TCM-4 tracking, the estimated P2 altitude was 138 km (down 72 km from the targeted value of 210 km). At or shortly after this became known, it was decided to undertake the building of TCM-5 for possible usage if the P2 altitude estimates were to continue to come down. A meeting was held with LMA on September 16, 1999, to agree upon a schedule. The plan was to build and test a maneuver having the same direction as TCM-4.

On September 16, 1999, the altitude estimate for P2 had moved back up to the region of 150 to 170 km. A TCM-5 delta-V magnitude was chosen that would raise P2 back up to the value of 210 km targeted at TCM-4, resulting in a periapsis altitude increase of roughly 44 km, if used. The date scheduled for TCM-5 was September 20, 1999.

A nominal MOI sequence had been sent to the spacecraft earlier as a precautionary measure in case the final upload could not be sent. On September 17, 1999, the project decided to delete the planned final MOI sequence load and to continue with the nominal sequence. Navigation was not aware of this decision until September 19, 1999.

The possibility of doing a P2 periapsis raise maneuver (PRM) by using a TCM at A1 (the first apoapsis of the established Mars orbit) was mentioned to project management by the LMA MCO Spacecraft Engineer in a discussion on September 17, 1999. In the September 17 to September 18 time period, orbit estimates of P2 were in the range of 141 to 172 km. Testing of the built TCM-5 sequence was completed on the night of September 17, 1999.

The TMC-5 maneuver sequence was run through the LMA STL concurrent with the MOI sequence "in the background," and no problems were encountered.

D. Board Meeting to Concur With No TCM-5 (9/18/99)

On Saturday morning September 18, 1999, at the request of the former MSP98 Program Manager (also a board member), a meeting was held to inform the Board of the possibility of executing a TCM-5 contingency maneuver and to obtain the Board's advice. Present at this meeting either in person or by phone were a limited number (3 of 8) of the Mars Surveyor Operations Standing Review Board members. Navigation personnel were not invited or among the attendees.

The Spacecraft Team had run the TCM-5 simulation in the STL on Friday night and felt that it would be safe to execute as proposed provided that the subsystem review was successfully completed. However, the project's position consisted of noting that the P2 altitude was still safe and that it would be risky to perform a TCM-5 during the same time period as a clocking-out MOI sequence. The tentative decision was made that there should be no TCM-5 unless the navigation solutions later in the day showed a significant problem. The Board chairman concurred with this. It was observed that an inefficient MOI sequence and a PRM could be used to raise P2 if warranted. However, at this point, the inefficient-MOI option had been ruled out, not only because LMA had never been formally requested to design this updated MOI sequence, but also because the nominal MOI sequence was already clocking out. This "clocking out" didn't make the inefficient MOI option infeasible, but it did entail some added complexity and risk compared to continuing with the sequence already in place.

The Navigation OD solution available at the September 18 afternoon status meeting did not indicate any significant change in the altitude estimates. The results of the morning meeting were not discussed at the afternoon status meeting.

At the scheduled TCM-5 Go/No-Go Meeting on Sunday September 19, 1999, the project announced that the no-go decision had been made the day before. Navigation accepted the decision not to perform TCM-5 after project management told Navigation on September 20, 1999, that the decision reflected the unanimous consensus of the Board.

E. TCM-5 Execution Window (9/20/99).

No TCM-5 was performed on September 20, 1999. Had TCM-5 been performed and raised the P1 altitude from its subsequently determined value of 57 km to perhaps 100 km, it is possible that the MCO spacecraft would have survived.

Appendix 6. Line-Project Relationships

The following is an excerpt from JPL D-12547, Revision 1, Soft Projectization.

Checks and Balances

The line organizations offer a wealth of knowledge, experience and data that cannot be replicated in each project element team and in fact should be leveraged across multiple projects to maximize resource-to-cost ratios. The Program Manager (PM) and the Project Element Managers (PEMs) are required to involve the line management in overseeing the project activities. The PM must involve line management in project status reviews to ensure quality communication, seek input, and build team spirit. The PEM must present the output of the element team (such as technical and resource plans, designs, and analyses) to the experts in the line organizations and seek their advice.

Experience has shown that the most effective mechanism for receiving these inputs is the "peer review." The peer review process, formally known as the Detailed Technical Review (DTR) process, is described in detail in the *JPL Standard for Review*, JPL D-10401. The PM should also require that the PEM seek technical/management peer review of a particular output, if it is felt that the extra scrutiny will improve the team's product.

As appropriate, line managers may suggest technical/peer reviews to the PEM. If they do not concur with the conclusions of the team or the actions of the PEM, they should provide their comments to the PM directly for response. The project should plan for and provide the needed resources for these check-and-balance functions.